mathematics

MDPI

Article

Continuous and Jump Variation in GARCH-MIDAS: Component Allocation and Volatility Forecasting

Mingxu Li [1] **and Sherry Zhefang Zhou** [1,2,*]

[1] Department of Mathematical Sciences, Beijing Normal-Hong Kong Baptist University, Zhuhai 519087, China; limingxu@bnbu.edu.cn
[2] Guangdong Provincial Key Laboratory of Interdisciplinary Research and Application for Data Science, Beijing Normal-Hong Kong Baptist University, Zhuhai 519087, China
* Correspondence: sherryzhou@bnbu.edu.cn

Abstract

This paper develops a generalized autoregressive conditional heteroskedasticity–mixed data sampling model with continuous and jump components (GARCH-MIDAS-CJ). It aligns the continuous–jump (CJ) decomposition of realized volatility with the long-run–short-run multiplicative structure. The model places block-smoothed daily continuous variation in the long-run mixed-data-sampling term and demeaned jump variation in the short-run Glosten–Jagannathan–Runkle GARCH (GJR-GARCH) recursion. Using high-frequency data for the Standard & Poor's 500 (S&P 500), we find a positive jump coefficient under both the Corsi–Pirino–Renò (CPR) and Andersen–Bollerslev–Dobrev (ABD) decompositions, supported by one-sided Wald, boundary-adjusted likelihood-ratio, and parametric-bootstrap tests. Relative to return- and realized-measure benchmarks, the forecasting value of the CJ allocation is competitive for multi-day and cumulative variance forecasts and in high-volatility periods, whereas established benchmarks can perform better for one-day and low-volatility forecasts. Under the CPR decomposition, the $h = 22$ average cumulative forecasts reduce the mean squared error (MSE) and QLIKE relative to the aggregate-realized-variance benchmark by 8.2% and 3.3%, respectively. One-day-ahead value-at-risk (VaR) results are strongest at the 1% and 0.5% tails; the 5% forecasts have exceedance rates above the nominal level. Across alternative information-timing conventions, forecast windows, and United States equity assets, the clearest gains remain concentrated at multi-day horizons, although their magnitude varies across settings. The findings support a persistence-based allocation of continuous and jump variation within the GARCH-MIDAS structure.

Keywords: realized volatility; continuous and jump variation; GARCH-MIDAS; volatility forecasting; value-at-risk

MSC: 62M10

Check for updates

Academic Editor: Zhifeng Dai

Received: 27 July 2026
Revised: 27 August 2026
Accepted: 28 August 2026
Published: 1 September 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Volatility forecasting is a central problem in empirical finance because conditional variance enters portfolio allocation, derivative pricing, and risk management, and a useful model must describe both volatility dynamics and the distribution of future returns [1]. Since the autoregressive conditional heteroskedasticity (ARCH) model of Engle [2] and the generalized autoregressive conditional heteroskedasticity (GARCH) model of Bollerslev [3], extensions such as exponential GARCH (EGARCH) [4], Glosten–Jagannathan–Runkle

GARCH (GJR-GARCH) [5], component GARCH [6], and fractionally integrated GARCH [7] have addressed asymmetry, multiple persistence scales, and long memory. Their parsimony and interpretability remain valuable, but specifications based only on daily returns do not use the information contained in intraday price variation.

Among the major advances beyond return-only models, the realized-volatility literature established that high-frequency returns provide informative ex post measures of latent variation [8–11]. Realized measures enter forecasting models through heterogeneous autoregressive (HAR), Realized GARCH, high-frequency-based volatility (HEAVY), and related specifications [12–14]. Measurement challenges accompany this informational gain: sampling too finely exposes realized variance to market microstructure noise, while discontinuous price movements make total realized variation an economically heterogeneous object [15,16]. Robust realized measures and continuous–jump (CJ) decompositions address these issues, but they do not by themselves determine where the resulting information should enter a conditional-variance model.

A complementary advance is the separation of persistent and transitory volatility within multiplicative component models. The generalized autoregressive conditional heteroskedasticity–mixed data sampling (GARCH-MIDAS) model decomposes conditional variance into a short-run GARCH term and a slowly moving long-run component [17]. The framework has been used to link long-run stock market volatility to macroeconomic and financial conditions [18,19]. At the same time, same-frequency formulations allow the long-run component to be driven by daily realized information [20]. Related spline and realized EGARCH-MIDAS models likewise show the value of separating baseline persistence from short-lived fluctuations [21,22]. This daily-frequency bridge makes it possible to align high-frequency volatility measurement with the long-run–short-run structure without forcing the MIDAS component to operate only at monthly or lower frequencies.

Realized volatility is commonly passed to GARCH-MIDAS models as a single object, although its continuous and jump components can differ in persistence and forecasting content [23–25]. We examine three questions: (i) whether continuous variation is more persistent than jump variation; (ii) whether lagged jump variation adds short-run information conditional on the continuous-based MIDAS component; and (iii) whether the proposed allocation improves forecasts relative to aggregate realized variance and alternative allocations. Component-level persistence diagnostics show that continuous variation is more persistent than jump variation in the aligned sample, thereby motivating the proposed long-run–short-run allocation.

The paper makes two related methodological contributions. First, it develops a cross-horizon allocation of continuous and jump variation within a same-frequency GARCH-MIDAS structure. Both inputs are obtained from the same realized measure: block-smoothed continuous variation determines the persistent MIDAS component, while demeaned jump variation enters the short-run GJR-GARCH recursion. The allocation links the continuous–jump decomposition to distinct long-run and short-run volatility channels according to the persistence of each component.

Second, the proposed specification provides a nested empirical design for distinguishing the incremental contribution of the short-run jump term from the forecasting value of the complete continuous–jump allocation. Imposing a zero jump coefficient yields the continuous-only GARCH-MIDAS specification while leaving the long-run input unchanged. This comparison isolates the contribution of jump variation to short-run dynamics. Comparing the complete model with GARCH-MIDAS based on aggregate realized variance then evaluates whether decomposing and reallocating the same realized measure improves forecasting performance. The resulting hierarchy separates inference on the jump-component coefficient from comparisons of predictive accuracy across non-nested models.

We apply the proposed method to the Standard & Poor's 500 (S&P 500) stock market index. Under both the Corsi–Pirino–Renò (CPR) [26] and Andersen–Bollerslev–Dobrev (ABD) [27] decompositions, the full-sample estimates support a positive short-run role for jump variation. In the out-of-sample (OOS) evaluation, the clearest loss reductions occur at multi-day horizons and during high-volatility periods, whereas the one-day and low-volatility results are more mixed. For one-day-ahead value-at-risk (VaR), the coverage tests do not reject the GM-CJ specifications at the 1% and 0.5% tails, whereas exact coverage is rejected at the 5% tail despite their low quantile losses. The multi-day forecasting pattern remains visible across alternative financial assets, OOS forecast windows, and information conventions, while the proposed split is retained in nearly all focused model confidence set (MCS) comparisons with alternative component allocations.

The remainder of the paper is organized as follows. Section 2 is the literature review. Section 3 proposes our GARCH-MIDAS-CJ model. Section 4 describes the data. Section 5 presents the empirical results, and Section 6 concludes.

2. Literature Review

The literature relevant to the proposed allocation can be organized by how volatility information is measured and assigned across persistence channels. Return-based ARCH and GARCH models represent conditional variance through dynamic recursions [2,3]. Subsequent specifications accommodate asymmetry, multiple persistence scales, and slow decay through EGARCH, GJR-GARCH, component GARCH, and fractionally integrated GARCH [4–7]. These models establish the main conditional-variance architecture, but their volatility information is extracted from daily returns.

High-frequency data provide more informative ex post measures of quadratic variation [8–11]. Models using realized measures differ in how this information is transmitted. HAR represents heterogeneous persistence through daily, weekly, and monthly realized-variance components [12]; Realized GARCH links returns, latent conditional variance, and realized measures through a measurement equation [13]; and HEAVY models use realized measures to update conditional-variance dynamics [14]. Extensions using semi-variances and jump-related measures further show that different parts of realized variation can have distinct forecasting roles [28]. Persistence is therefore represented through lag aggregation, latent-state dynamics, or joint updating rather than through a common allocation rule.

A second mechanism separates persistent and transitory conditional variance. Spline-GARCH and GARCH-MIDAS use multiplicative components to distinguish a slowly moving volatility level from short-run dynamics [17,21,29]. Macroeconomic and financial variables can explain the persistent component [18,19], while same-frequency formulations allow daily realized information to enter the MIDAS term [20,22]. Recent extensions broaden the channels through which daily covariates, multiple frequencies, current information, regime changes, and short-run forecast errors affect the two components [30–34]. This literature provides the long-run–short-run architecture within which the allocation of decomposed realized information can be examined.

A third mechanism separates quadratic variation into continuous and jump components. Bipower variation and related jump-robust estimators identify continuous variation in the presence of discontinuous price movements [35,36], while subsequent procedures address local jumps, market microstructure noise, and finite-sample distortions [37–39]. Empirical evidence shows that jumps are episodic but economically meaningful and can differ from continuous variation in persistence and forecasting content [23–25]. Their contribution also varies with the jump measure, forecast horizon, market conditions, and the way decomposed information enters the model [28,40–43].

Table 1 summarizes how the closest realized-measure frameworks represent persistence and transmit realized information to conditional variance.

Table 1. Mechanism comparison of realized-measure volatility frameworks.

Framework	Persistence Representation	Continuous and Jump Information	Link to Conditional Variance
GARCH-MIDAS/GM-RV [17,20]	A weighted MIDAS history determines a persistent component.	Aggregate realized variance may drive the MIDAS term without an explicit continuous–jump allocation.	The persistent MIDAS component multiplies the daily GARCH component.
HAR-CJ [24]	Daily, weekly, and monthly lags represent heterogeneous persistence.	Continuous and jump measures enter as separate observed predictors.	The components forecast realized variance through an additive autoregressive structure.
Realized GARCH [13,28]	Latent conditional variance follows a persistent dynamic process.	Realized and decomposed measures enter through the measurement system.	Returns, latent variance, and realized measures are estimated jointly.
HEAVY [14]	Realized measures update the dynamics of conditional variance.	Realized inputs need not be assigned to separate continuous and jump horizons.	Return and realized-measure dynamics evolve through linked equations.
GARCH-MIDAS-CJ	Block-smoothed continuous variation determines the persistent MIDAS component.	Centered jump variation enters the short-run GJR-GARCH recursion.	The two inputs from the same realized measure enter distinct components of $\sigma_t^2 = \tau_t g_t$.

Notes: GM-RV denotes GARCH-MIDAS with aggregate realized variance. HAR-CJ denotes a heterogeneous autoregressive specification with continuous and jump components. HEAVY denotes high-frequency-based volatility.

The proposed model uses the multiplicative GARCH-MIDAS structure to allocate the continuous and jump parts of the same realized measure across persistent and short-run volatility channels. Persistence diagnostics evaluate the premise of this allocation, the nested GM-C–GM-CJ comparison identifies the incremental short-run jump contribution, and comparisons with GM-RV and alternative allocations assess its forecasting implications.

3. Materials and Methods

3.1. Realized Variance and Continuous–Jump Decomposition

Let $p_{t,i}$ denote the logarithmic price observed on day t at intraday time $i = 0, 1, \ldots, M$, where M is the number of intraday intervals. The corresponding intraday return $r_{t,i}$ is

$$r_{t,i} = p_{t,i} - p_{t,i-1}, \tag{1}$$

and daily realized variance is defined as

$$RV_t = \sum_{i=1}^{M} r_{t,i}^2. \tag{2}$$

As the sampling interval shrinks, RV_t converges to the quadratic variation QV_t of the log-price process. In the presence of jumps, quadratic variation can be decomposed as

$$QV_t = IV_t + JV_t, \tag{3}$$

where IV_t is integrated continuous variation and JV_t is jump variation.

We consider two complementary jump-detection procedures. The CPR test of Corsi et al. [26] uses daily aggregates of intraday returns and detects jumps from the discrepancy between realized variance and an estimate of continuous variation. The ABD test of Andersen et al. [24] examines standardized returns interval by interval and can locate multiple jumps within the same trading day. CPR and ABD therefore represent, respectively, the daily aggregate and intraday approaches to jump detection. The CPR decomposition is used in the main analysis, while the ABD decomposition provides an alternative based on a different level of information aggregation.

The CPR procedure builds on the bipower-variation framework of Barndorff-Nielsen and Shephard [35,36]. Define

$$\mu_a = E(|Z|^a) = \frac{2^{a/2}\Gamma((a+1)/2)}{\sqrt{\pi}}, \qquad a > 0, \qquad Z \sim N(0,1), \tag{4}$$

where a is the order of the absolute moment. Bipower variation is given by

$$BV_t = \mu_1^{-2} \sum_{i=2}^{M} |r_{t,i}||r_{t,i-1}|. \tag{5}$$

While realized variance contains both continuous and jump variation, BV_t consistently estimates the continuous component under finite-activity jumps. The CPR test of Corsi et al. [26] improves robustness to jumps by constructing threshold-corrected bipower and tripower measures. For each intraday return, the threshold is

$$\vartheta_{t,i} = c_\vartheta^2 \widehat{V}_{t,i}, \tag{6}$$

where $\widehat{V}_{t,i}$ is a local variance estimate. Corsi et al. [26] shows that the estimator becomes relatively insensitive to the scale parameter when $c_\vartheta \geq 3$; accordingly, we set $c_\vartheta = 3$.

Let $Z_a(x,y)$ denote the corrected threshold transformation:

$$Z_a(x,y) = \begin{cases} |x|^a, & x^2 \leq y, \\ E\left(|X|^a \mid X^2 > y\right), & x^2 > y, \quad X \sim N\left(0, y/c_\vartheta^2\right). \end{cases} \tag{7}$$

Thus, observations below the threshold enter in their original form, whereas larger observations are replaced by their conditional Gaussian moments. The power a is determined by the realized measure being constructed: $a = 1$ for corrected threshold bipower variation and $a = 4/3$ for corrected threshold tripower variation.

The two corrected measures are

$$CTBPV_t = \mu_1^{-2} \sum_{i=2}^{M} Z_1(r_{t,i}, \vartheta_{t,i}) Z_1(r_{t,i-1}, \vartheta_{t,i-1}) \tag{8}$$

and

$$CTTriPV_t = M\mu_{4/3}^{-3} \sum_{i=3}^{M} Z_{4/3}(r_{t,i}, \vartheta_{t,i}) Z_{4/3}(r_{t,i-1}, \vartheta_{t,i-1}) Z_{4/3}(r_{t,i-2}, \vartheta_{t,i-2}). \tag{9}$$

Thresholding limits the influence of discontinuous observations on the estimated continuous component, while the conditional-moment correction avoids mechanically discarding all observations above the threshold. The CPR statistic is

$$Z_t^{\mathrm{CPR}} = \frac{1 - CTBPV_t / RV_t}{\sqrt{\left[\left(\frac{\pi}{2}\right)^2 + \pi - 5\right] M^{-1} \max\left\{1, \frac{CTTriPV_t}{CTBPV_t^2}\right\}}}. \tag{10}$$

Under the null hypothesis of no jumps, Z_t^{CPR} is asymptotically standard normal. Given a daily significance level α_J, the CPR decomposition is

$$J_t^{\mathrm{CPR}} = \mathbf{1}\left\{ Z_t^{\mathrm{CPR}} > \Phi_{1-\alpha_J}^{-1} \right\} \max\{RV_t - CTBPV_t, 0\}, \qquad C_t^{\mathrm{CPR}} = RV_t - J_t^{\mathrm{CPR}}. \tag{11}$$

Here, $\mathbf{1}\{A\}$ equals one if event A occurs and zero otherwise, and Φ_q^{-1} denotes the q-th quantile of the standard normal distribution.

The ABD test of Andersen et al. [27] identifies jumps directly from standardized intraday returns. Its interval-specific statistic is

$$T_{t,i}^{\mathrm{ABD}} = \frac{r_{t,i}}{\sqrt{M^{-1}BV_t}}. \tag{12}$$

Under the null of no jumps and locally constant intraday volatility, $T_{t,i}^{\mathrm{ABD}}$ is approximately standard normal. Because the test is applied to each of the M intraday returns, the interval-level significance level is adjusted as

$$\alpha_J^* = 1 - \left(1 - \alpha_J\right)^{1/M}. \tag{13}$$

An intraday return is classified as a jump when

$$\left| T_{t,i}^{\mathrm{ABD}} \right| > \Phi_{1-\alpha_J^*/2}^{-1}. \tag{14}$$

Let

$$\mathcal{J}_t^{\mathrm{ABD}} = \left\{ i : \left| T_{t,i}^{\mathrm{ABD}} \right| > \Phi_{1-\alpha_J^*/2}^{-1} \right\}, \qquad \text{and} \qquad N_t^{\mathrm{ABD}} = \left| \mathcal{J}_t^{\mathrm{ABD}} \right| \tag{15}$$

denote the detected jump intervals and their number, respectively. Following the variance-allocation approach of Audrino and Hu [44], the average variation associated with a non-jump return is estimated by

$$\widehat{v}_t^{\,c} = \frac{1}{M - N_t^{\mathrm{ABD}}} \sum_{i \notin \mathcal{J}_t^{\mathrm{ABD}}} r_{t,i}^2. \tag{16}$$

Under the ABD decision rule in Equation (14), $N_t^{\mathrm{ABD}} < M$ for the daily significance level $\alpha_J = 5\%$ used in this study. Hence, $\widehat{v}_t^{\,c}$ is well defined; a proof is provided in Appendix A. The ABD-based components are then

$$J_t^{\mathrm{ABD}} = \sum_{i \in \mathcal{J}_t^{\mathrm{ABD}}} \left(r_{t,i}^2 - \widehat{v}_t^{\,c} \right), \qquad C_t^{\mathrm{ABD}} = RV_t - J_t^{\mathrm{ABD}}, \tag{17}$$

with $J_t^{\mathrm{ABD}} = 0$ when no intraday jump is detected. Here, $C_t^{(d)}$ and $J_t^{(d)}$ denote the empirical measures of continuous and jump variation constructed by procedure $d \in \{\mathrm{CPR}, \mathrm{ABD}\}$. Throughout the remainder of the paper, the CPR or ABD superscript identifies the empirical decomposition procedure. Once a decomposition is selected for a specification, the procedure superscript is suppressed, so that C_t and J_t denote the corresponding empirical continuous and jump series. Parameters and model states are written without hats, while parameter estimates and fitted state paths are written with hats.

3.2. Benchmark GARCH-MIDAS Structure

The starting point is the multiplicative component GARCH-MIDAS framework of Engle et al. [17]. The model separates conditional variance into a short-run component and a long-run component. In the conventional mixed-frequency setting, daily returns are

indexed by i within a lower-frequency period t, where t may represent a month or a quarter. Let $r_{i,t}$ denote the daily return on day i of period t. The return equation is

$$r_{i,t} = \mu + \varepsilon_{i,t}, \qquad \varepsilon_{i,t} = \sqrt{g_{i,t}\tau_t}\, z_{i,t}, \qquad z_{i,t} \sim i.i.d.(0,1), \tag{18}$$

Here, μ is the conditional mean, $\varepsilon_{i,t}$ is the return innovation, $z_{i,t}$ is a standardized innovation, and $\mathcal{F}_{i-1,t}$ denotes the information available before day i of period t. The corresponding conditional variance is

$$\sigma^2_{i,t} = \mathrm{Var}(r_{i,t} \mid \mathcal{F}_{i-1,t}) = g_{i,t}\tau_t. \tag{19}$$

Here, $g_{i,t}$ is the short-run volatility component, while τ_t is the long-run volatility component. The defining feature of the model is that $g_{i,t}$ changes at the daily frequency, whereas τ_t changes only at the lower frequency. This structure allows the model to capture both high-frequency volatility clustering and slow-moving changes in the baseline volatility level. The multiplicative decomposition is closely related to the component-volatility literature and provides a flexible alternative to purely additive component models [6,20,21].

The short-run component is usually specified as a unit-mean GARCH-type process. Following the leverage specification of Glosten et al. [5], and consistent with the GARCH-MIDAS applications in Engle et al. [17] and Conrad and Kleen [20], we use a GJR-GARCH form:

$$g_{i,t} = \left(1 - \alpha - \frac{\gamma}{2} - \beta\right) + \left(\alpha + \gamma \mathbf{1}_{\{\varepsilon_{i-1,t}<0\}}\right)\frac{\varepsilon^2_{i-1,t}}{\tau_t} + \beta g_{i-1,t}. \tag{20}$$

Here, α is the symmetric news coefficient, γ is the additional response to a negative return shock, and β is the autoregressive coefficient of the short-run component. The indicator $\mathbf{1}_{\{\varepsilon_{i-1,t}<0\}}$ equals one for a negative lagged return innovation and zero otherwise. The normalization $1 - \alpha - \gamma/2 - \beta$ implies that the short-run component fluctuates around one, provided that the innovation is symmetric and the persistence condition

$$\alpha + \frac{\gamma}{2} + \beta < 1 \tag{21}$$

holds. Under this normalization, τ_t can be interpreted as the long-run volatility level around which the short-run component moves.

In the original mixed-frequency GARCH-MIDAS model, the long-run component is driven by a lower-frequency explanatory variable, such as a monthly or quarterly macroeconomic variable. Let X_t denote the low-frequency variable observed at period t. The long-run component is commonly written as

$$\tau_t = \exp\left(m + \theta \sum_{k=1}^{K} \phi_k(w_1, w_2) X_{t-k}\right), \tag{22}$$

where m is the long-run intercept, θ is the loading on the low-frequency variable, K is the number of low-frequency lags, and $\phi_k(w_1, w_2)$ is the weight assigned to lag k. The parameters w_1 and w_2 determine the shape of the weighting function. The MIDAS weights are parameterized by the two-parameter Beta function:

$$\phi_k(w_1, w_2) = \frac{\left(\frac{k}{K+1}\right)^{w_1-1}\left(1 - \frac{k}{K+1}\right)^{w_2-1}}{\sum_{j=1}^{K}\left(\frac{j}{K+1}\right)^{w_1-1}\left(1 - \frac{j}{K+1}\right)^{w_2-1}}, \qquad k = 1, \ldots, K. \tag{23}$$

The normalization ensures that the weights sum to one. The two shape parameters allow for decreasing, increasing, or hump-shaped lag profiles, thereby summarizing a potentially long history of low-frequency information with only a small number of parameters [17,29]. The exponential specification in Equation (22) ensures that the long-run component remains positive.

Although the original GARCH-MIDAS model was designed for mixed-frequency applications, the same multiplicative structure can also be used when the long-run driver is observed at the daily frequency. Conrad and Kleen [20] discuss the multiplicative GARCH-MIDAS model as a general component-volatility framework and note that when the period contains a single observation, the within-period index can be dropped and both components may be treated at the same frequency. This observation is important for realized-measure applications, because realized volatility and its continuous and jump components are naturally available at the daily frequency. In this same-frequency version, the model can be written as

$$r_t = \mu + \varepsilon_t, \qquad \varepsilon_t = \sqrt{g_t \tau_t}\, z_t, \qquad \sigma_t^2 = g_t \tau_t, \tag{24}$$

where g_t and τ_t are both indexed by the trading day t. The short-run recursion becomes

$$g_t = \left(1 - \alpha - \frac{\gamma}{2} - \beta\right) + \left(\alpha + \gamma \mathbf{1}_{\{\varepsilon_{t-1}<0\}}\right) \frac{\varepsilon_{t-1}^2}{\tau_{t-1}} + \beta g_{t-1}. \tag{25}$$

This daily version provides the benchmark structure for the present paper. It keeps the economic interpretation of GARCH-MIDAS—a short-run component fluctuating around a persistent long-run component—while allowing the long-run component to be driven by daily realized measures. This is also consistent with the broader realized-measure literature, which shows that high-frequency realized measures contain substantially more information about latent volatility than squared daily returns [8,11,13].

In the empirical implementation, the innovation z_t is assumed to follow a standardized Student-t distribution with degrees of freedom $\nu > 2$. The Student-t specification is widely used in GARCH-type volatility modeling because equity returns typically exhibit excess kurtosis relative to the Gaussian distribution [3,17]. Thus,

$$z_t \sim t_\nu^{std}(0,1), \tag{26}$$

where $t_\nu^{std}(0,1)$ denotes the Student-t distribution standardized to have mean zero and variance one.

3.3. The Proposed GARCH-MIDAS-CJ Model

The proposed GARCH-MIDAS-CJ model links the continuous–jump decomposition of realized variance to the long-run–short-run decomposition of conditional variance. Continuous variation is used to characterize the persistent volatility level, whereas jump variation enters the short-run recursion as a measure of discontinuous market shocks. For a selected decomposition $d \in \{\mathrm{CPR}, \mathrm{ABD}\}$, we suppress the procedure superscript after defining the empirical components. The connection between the two decompositions is then summarized by

$$\begin{aligned} RV_s &= C_s + J_s, \\ \{C_s\}_{s<t} &\xrightarrow{\text{block aggregation and MIDAS weighting}} \tau_t, \\ J_{t-1} &\longmapsto J_{t-1} - \bar{J} \longmapsto g_t, \\ (\tau_t, g_t) &\longmapsto \sigma_t^2 = \tau_t g_t. \end{aligned} \tag{27}$$

Equation (27) defines the persistence-based specification examined here; alternative placements of the two realized components are evaluated in Section 5.4.

A daily realized measure can contain movements at the same horizon as those modeled by the short-run component. Feeding each daily observation directly into the MIDAS filter could therefore make the long-run component respond too strongly to day-to-day variation. Conrad and Kleen [20] note that, even when the explanatory variable is observed daily, the long-run component should remain a smooth function of a sufficiently long history of that variable. More directly, Borup and Jakobsen [22] separate the effects of realized measures by assigning their most immediate variation primarily to the short-run component and using older, aggregated information to determine the long-run component. We adopt their non-overlapping aggregation device for the continuous component.

The long-run component is specified as

$$\tau_t = \exp\left(m + \lambda \sum_{k=1}^{K} \phi_k(w_2) \bar{c}_{t-\ell,k}^{(N)} \right), \tag{28}$$

where λ is the loading on block-averaged log continuous variation, K is the number of non-overlapping blocks, and ℓ determines the most recent observation entering the MIDAS filter. The k-th non-overlapping block average is

$$\bar{c}_{t-\ell,k}^{(N)} = \frac{1}{N} \sum_{j=1}^{N} \log C_{t-\ell-N(k-1)-j+1}, \qquad k = 1, \ldots, K, \tag{29}$$

where N is the number of daily observations in each block, and C_t denotes the empirical daily continuous-variation measure supplied by the selected CPR or ABD decomposition in Equations (11) and (17), respectively. Across the empirical samples, no zero, negative, or non-finite estimates of continuous variation occur. The minimum CPR and ABD estimates are 0.00302 and 0.00229 for the S&P 500, and 0.00935 and 0.00883 for the NASDAQ 100, respectively. Hence, $\log C_t$ is defined for every observation entering the MIDAS component. Thus, each block contains N consecutive observations of log continuous variation, and the K blocks collectively cover NK trading days without overlap. The empirical implementation uses $N = 22$ observations per block, approximately one trading month (this block length follows the monthly convention used in related daily-frequency GARCH-MIDAS specifications and is held fixed across models to preserve comparability [20,22]), and $K = 24$ blocks, giving the MIDAS filter a maximum support of 528 trading days, or approximately two trading years. Its effective memory depends on the concentration of the estimated weights. The sensitivity to shorter and longer monthly spans and to weekly blocks is reported in Appendix D.

The information-timing parameter ℓ determines the most recent observation entering the MIDAS filter. Under the convention $\ell = 2$ in Borup and Jakobsen [22], the most recent block ends at $t - 2$, leaving information dated $t - 1$ to the short-run recursion. The alternative specification sets $\ell = 1$, so that the most recent block ends at $t - 1$; this alternative is used as a timing robustness check.

Following the restricted Beta-weighting specification used by Conrad et al. [45], based on Engle et al. [17], we impose $w_1 = 1$, so that

$$\phi_k(w_2) \equiv \phi_k(1, w_2) = \frac{\left(1 - \frac{k}{K+1}\right)^{w_2 - 1}}{\sum_{j=1}^{K} \left(1 - \frac{j}{K+1}\right)^{w_2 - 1}}, \qquad k = 1, \ldots, K. \tag{30}$$

This restricted scheme leaves the decay parameter w_2 to be estimated while reducing the lag profile to a single shape parameter. For $w_2 > 1$, the weights are positive, sum to one, and decline monotonically with the lag order, with larger values of w_2 producing faster decay. The unrestricted two-parameter Beta scheme additionally permits hump-shaped profiles. We use the restricted form to retain a parsimonious declining MIDAS filter and limit weak identification between nearly equivalent two-parameter weighting profiles [20,22].

Following the GARCH-MIDAS-X construction of Amendola et al. [30], we augment the short-run GJR-GARCH recursion with a demeaned nonnegative variable observed at the same frequency as daily returns. Their specification introduces the additional variable X_t in the population-centered form $z[X_{t-1} - E(X_{t-1})]$. Under standardized symmetric innovations and a stationary solution with finite first moments, the centered term has zero expectation and therefore does not alter the population normalization $E(g_t) = 1$. This separates the baseline scale of g_t from movements in the additional variable. We adopt this device but use lagged jump variation as the short-run variable:

$$g_t = \left(1 - \alpha - \frac{\gamma}{2} - \beta\right) + \left(\alpha + \gamma \mathbf{1}_{\{\varepsilon_{t-1}<0\}}\right) \frac{\varepsilon_{t-1}^2}{\tau_{t-1}} + \beta g_{t-1} + \zeta(J_{t-1} - \bar{J}). \tag{31}$$

The centering term $\bar{J}$ is computed over the observations entering the likelihood:

$$\bar{J} = \frac{1}{T_0} \sum_{t \in \mathcal{T}} J_{t-1}, \qquad T_0 = |\mathcal{T}|, \tag{32}$$

where $\mathcal{T}$ denotes the aligned estimation sample.

The coefficient $\zeta \geq 0$ measures the response of short-run volatility to jump variation relative to its usual level. The sign restriction imposes a nonnegative response to above-average lagged jump variation, with $\zeta = 0$ defining the nested benchmark of no incremental jump effect. Whereas Amendola et al. [30] formulate the short-run extension for a general same-frequency volatility determinant, our specification uses the jump component obtained from the same realized-variance decomposition that supplies continuous variation to the long-run MIDAS term. By construction, $J_{t-1} - \bar{J}$ has a sample mean of zero over the likelihood window. This is the sample analogue of population centering by $E(J_t)$ used in Amendola et al. [30]. Under a stationary solution with finite first moments and standardized symmetric innovations independent of past information, population centering yields the normalization $E(g_t) = 1$.

In finite samples, centering by $\bar{J}$ removes the mechanical contribution of the average jump regressor to the baseline of g_t, but it does not impose that the recursively generated g_t path has sample mean exactly equal to one. Initialization and the realized shock sequence may therefore produce a finite-sample mean different from one. If J_t is stationary and ergodic with $E|J_t| < \infty$, then $\bar{J} \to E(J_t)$ as the estimation window expands. Whenever the model is re-estimated in the rolling exercise, $\bar{J}$ is recomputed using only the aligned observations in the current estimation window. Accordingly, ζ measures the short-run response to jump variation relative to its estimation-window mean.

Figure 1 summarizes the information allocation in the proposed model. The CPR or ABD decomposition supplies both realized-measure inputs: block-aggregated continuous variation determines the long-run MIDAS state, while centered jump variation and the lagged daily return innovation update the short-run GJR-GARCH state. Their product gives the conditional variance used for estimation and forecasting.

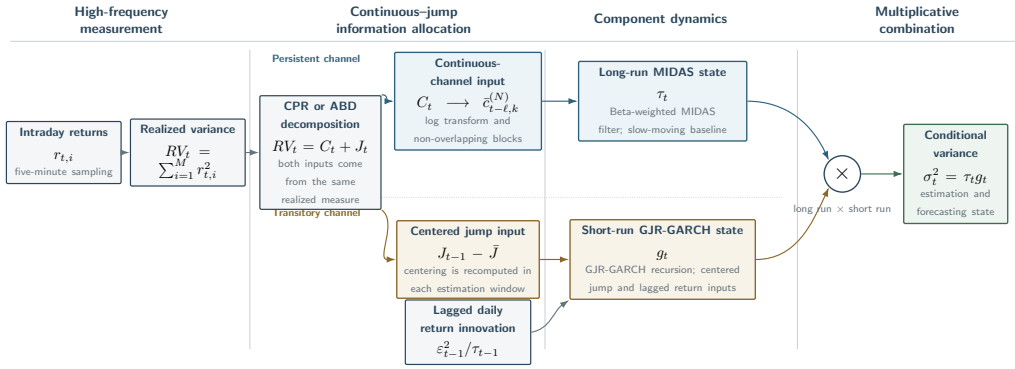


Figure 1. Information allocation in the GARCH-MIDAS-CJ model. Notes: CPR and ABD denote the Corsi–Pirino–Renò and Andersen–Bollerslev–Dobrev decompositions, respectively. The continuous and jump inputs are derived from the same realized-variance measure. The upper branch shows the persistent channel from block-aggregated continuous variation to the long-run MIDAS state τ_t. The lower branch shows how centered jump variation and the lagged daily return innovation enter the short-run state g_t. The two states combine multiplicatively as $\sigma_t^2 = \tau_t g_t$.

The positivity implication of the centered jump term can be stated directly. Let $\mu_J = E(J_t)$ denote the population mean of jump variation and define

$$c = 1 - \alpha - \frac{\gamma}{2} - \beta, \qquad A_{t-1} = \beta + \left(\alpha + \gamma \mathbf{1}_{\{\varepsilon_{t-1}<0\}}\right) z_{t-1}^2. \tag{33}$$

Using $\varepsilon_{t-1}^2/\tau_{t-1} = g_{t-1} z_{t-1}^2$, the population-centered counterpart of Equation (31) is

$$g_t = A_{t-1} g_{t-1} + c + \zeta\left(J_{t-1} - \mu_J\right). \tag{34}$$

Because $A_{t-1} \geq 0$, $J_{t-1} \geq 0$, and $\zeta \geq 0$, a simple sufficient condition for strict positivity is

$$c - \zeta \mu_J > 0, \qquad g_0 > 0. \tag{35}$$

Under this condition, Equation (34) implies $g_t > 0$ by induction. In the sample-centered implementation, $\bar{J}$ replaces μ_J, giving the corresponding sample condition $c - \zeta \bar{J} > 0$. Imposing $\zeta = 0$ removes this channel, yielding the continuous-only GARCH-MIDAS (GM-C) model:

$$H_0 : \zeta = 0 \qquad \text{against} \qquad H_1 : \zeta > 0. \tag{36}$$

For the augmented short-run recursion, we maintain that the sequence $\{(A_t, c + \zeta(J_t - \mu_J))\}$ is jointly stationary and ergodic and satisfies

$$E[\log A_t] < 0, \qquad E\left[\log^+ \left|c + \zeta(J_t - \mu_J)\right|\right] < \infty, \tag{37}$$

where $\log^+(x) := \max\{\log(x), 0\}$. Together with the positivity condition above, these conditions provide a sufficient basis for a unique causal, strictly stationary, and ergodic solution for g_t [46,47]. We also maintain that the realized inputs determining τ_t and J_t are jointly stationary and ergodic. Existence of a moment of order p requires the corresponding order-p stochastic-recurrence contraction and finite input moments, while moments of the complete conditional variance require $E[(\tau_t g_t)^p] < \infty$.

Under conditional symmetry and the unit-variance normalization, $E(A_t) = \alpha + \gamma/2 + \beta$. Hence, $\alpha + \gamma/2 + \beta < 1$ implies the logarithmic contraction condition whenever the expectation is well defined. This conventional GJR-GARCH restriction is not sufficient in isolation for the complete augmented process because it does not ensure the stationarity, integrability,

or required moments of the jump and long-run inputs. The restriction $\nu > 2$ ensures a finite innovation variance; estimation and robust inference additionally rely on identification, differentiability, finite score moments, and nonsingular information matrices.

3.4. Inference on the Jump-Component Coefficient

Let $\boldsymbol{\psi}$ collect the model parameters and define the standardized residual as

$$u_t = \frac{r_t - \mu}{\sqrt{g_t \tau_t}}. \tag{38}$$

Under standardized Student-t innovations, the contribution of observation t to the log-likelihood is

$$\begin{aligned} \ell_t(\boldsymbol{\psi}) = \log\Gamma\left(\frac{\nu+1}{2}\right) - \log\Gamma\left(\frac{\nu}{2}\right) - \frac{1}{2}\log\{(\nu-2)\pi\} \\ - \frac{1}{2}\log(g_t\tau_t) - \frac{\nu+1}{2}\log\left(1 + \frac{u_t^2}{\nu-2}\right). \end{aligned} \tag{39}$$

The estimator maximizes

$$L(\boldsymbol{\psi}) = \sum_{t=t_0}^{T} \ell_t(\boldsymbol{\psi}). \tag{40}$$

The starting date t_0 is determined by the MIDAS initialization requirement. Constructing K non-overlapping blocks of length N requires NK lagged daily observations before the first likelihood contribution can be evaluated.

The parameter space imposes

$$\alpha \geq 0, \qquad \gamma \geq 0, \qquad \beta \geq 0, \qquad \alpha + \frac{\gamma}{2} + \beta < 1, \qquad \nu > 2, \qquad \zeta \geq 0. \tag{41}$$

During likelihood evaluation, no dynamic floor or post-recursion truncation is applied. Candidate parameter values that generate a non-positive or non-finite path for g_t, τ_t, or $\sigma_t^2 = g_t\tau_t$ are treated as infeasible. This pathwise numerical rule is distinct from the analytical sufficient condition in Equation (35).

Because ζ is restricted to be nonnegative, the null value $\zeta = 0$ lies on the boundary of the admissible parameter space. Conventional two-sided Wald inference and the usual chi-squared likelihood-ratio approximation are therefore inappropriate. We assess

$$H_0 : \zeta = 0 \qquad \text{against} \qquad H_1 : \zeta > 0 \tag{42}$$

using a one-sided Wald calculation, a boundary-adjusted likelihood-ratio test, and a conditional parametric-bootstrap likelihood-ratio test.

For the Wald calculation, let

$$s_t(\boldsymbol{\psi}) := \frac{\partial \ell_t(\boldsymbol{\psi})}{\partial \boldsymbol{\psi}}, \qquad \bar{s} := \frac{1}{T_0}\sum_{t\in\mathcal{T}} s_t(\widehat{\boldsymbol{\psi}}_U), \tag{43}$$

where $\widehat{\boldsymbol{\psi}}_U$ is the constrained maximum-likelihood estimate of the unrestricted GARCH-MIDAS-CJ (GM-CJ) model. Define

$$\widehat{H}_U := -\frac{1}{T_0}\sum_{t\in\mathcal{T}} \frac{\partial^2 \ell_t(\widehat{\boldsymbol{\psi}}_U)}{\partial \boldsymbol{\psi}\, \partial \boldsymbol{\psi}'}, \tag{44}$$

$$\widehat{S}_U := \frac{1}{T_0}\sum_{t\in\mathcal{T}} \left[s_t(\widehat{\boldsymbol{\psi}}_U) - \bar{s}\right]\left[s_t(\widehat{\boldsymbol{\psi}}_U) - \bar{s}\right]', \tag{45}$$

$$\widehat{V}_U := \frac{1}{T_0}\widehat{H}_U^{-1}\widehat{S}_U\widehat{H}_U^{-1}. \tag{46}$$

This is the Bollerslev–Wooldridge-type quasi-maximum-likelihood (QML)-robust sandwich covariance estimator [48]. The standard error used for the jump coefficient is $\widehat{\mathrm{se}}(\widehat{\zeta}) = \sqrt{[\widehat{V}_U]_{\zeta\zeta}}$, where $[\widehat{V}_U]_{\zeta\zeta}$ denotes the diagonal element corresponding to ζ. The one-sided Wald statistic and p-value are

$$W_\zeta := \frac{\widehat{\zeta}}{\widehat{\mathrm{se}}(\widehat{\zeta})}, \qquad p_W := 1 - \Phi(W_\zeta), \tag{47}$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function. Because the null is located on the boundary, this Wald result is reported as supplementary rather than primary inference.

The likelihood-ratio statistic compares the unrestricted GM-CJ model with the restricted GM-C model obtained by imposing $\zeta = 0$:

$$LR_\zeta := 2\big[L(\widehat{\boldsymbol{\psi}}_U) - L(\widehat{\boldsymbol{\psi}}_R)\big], \tag{48}$$

where $\widehat{\boldsymbol{\psi}}_R$ denotes the restricted maximum-likelihood estimate. Under the regular single-boundary-parameter setting, the null distribution is

$$LR_\zeta \xrightarrow{d} \frac{1}{2}\chi_0^2 + \frac{1}{2}\chi_1^2, \tag{49}$$

where χ_0^2 is degenerate at zero and χ_1^2 is a chi-squared distribution with one degree of freedom [49,50]. For an observed $LR_\zeta > 0$, the corresponding boundary-mixture p-value is

$$p_{\mathrm{mix}} := \frac{1}{2}\Big[1 - F_{\chi_1^2}(LR_\zeta)\Big], \tag{50}$$

where $F_{\chi_1^2}(\cdot)$ denotes the cumulative distribution function of χ_1^2. This calculation is referred to as the boundary-adjusted likelihood-ratio test below.

The 50:50 mixture approximation in Equation (49) applies when ζ is the only active boundary parameter and the remaining nuisance parameters are locally identifiable and interior under the null. In the present estimates, α is also on its boundary, so the null distribution may depend on the joint geometry of the parameter constraints. Primary finite-sample inference therefore uses the conditional parametric bootstrap, with the mixture p-value reported as a single-boundary benchmark. Returns are simulated under the restricted GM-C model with $\zeta = 0$, using the restricted parameter estimates and standardized Student-t innovations while holding fixed the observed realized-measure sequences entering the long-run and jump specifications. For each bootstrap sample b, both GM-C and GM-CJ are re-estimated under the original parameter constraints, and

$$LR_\zeta^{(b)} := 2\Big[L^{(b)}(\widehat{\boldsymbol{\psi}}_U^{(b)}) - L^{(b)}(\widehat{\boldsymbol{\psi}}_R^{(b)})\Big] \tag{51}$$

is calculated. Here, $L^{(b)}$ is the log-likelihood for bootstrap sample b, and $\widehat{\boldsymbol{\psi}}_U^{(b)}$ and $\widehat{\boldsymbol{\psi}}_R^{(b)}$ are its unrestricted and restricted estimates. The bootstrap p-value is

$$p_{\mathrm{boot}} := \frac{1 + \sum_{b=1}^{B} \mathbf{1}\Big\{LR_\zeta^{(b)} \geq LR_\zeta\Big\}}{B+1}, \tag{52}$$

where B is the number of successful bootstrap replications and $\mathbf{1}\{\cdot\}$ equals one when its argument is true and zero otherwise. Re-estimating both models in every replication allows the bootstrap distribution to reflect the same inequality constraints and boundary solutions

encountered in the original sample. A bootstrap replication is classified as unsuccessful if either the restricted or unrestricted fit fails to return a converged, finite, and constraint-feasible likelihood. The failure reason is recorded, the unsuccessful draw is not replaced, and it is excluded from both B and the tail-probability calculation. All requested replications reported below were successful, so this rule does not affect the resulting p-values.

3.5. Computational Implementation

The empirical analysis was implemented in R 4.5.2 under Windows 11. The volatility recursions were compiled with Rcpp 1.1.2, and the implementation used data.table 1.18.4, numDeriv 2016.8-1.1, forecast 9.0.2, and rugarch 1.5-6 for data processing, numerical derivatives, auxiliary time-series estimation, Realized GARCH estimation, and MCS evaluation. Constrained quasi-maximum-likelihood estimation uses bounded L-BFGS-B optimization with a maximum of 5000 iterations and `factr` $= 10^7$, followed by `nlminb` refinement. As an additional numerical check, the estimates are verified using `constrOptim` under the same parameter constraints. A candidate solution is retained only when the solver reports convergence, the parameters and objective are finite, all parameter constraints are satisfied, and the resulting g_t, τ_t, and σ_t^2 paths are finite and strictly positive.

Full-sample estimation uses 20 initial values. Rolling estimation begins from the preceding window's estimate and, when needed, proceeds to deterministic and pseudo-random starts, with up to five initial values. Each restricted and unrestricted fit in the jump-coefficient bootstrap uses three initial values. The CPR and ABD bootstrap tests use base seeds 20260623 and 20260624, respectively, and produce 499 and 999 successful replications. The 90% MCS uses 5000 stationary-bootstrap replications, with the block length selected from the dependence in the pairwise loss-difference series.

4. Data

The empirical analysis uses daily and five-minute intraday data for two liquid United States equity indices: the S&P 500 and the NASDAQ 100. The S&P 500 serves as the main empirical application, while the NASDAQ 100 is used for the cross-index robustness analysis. The data are obtained from the Wind Economic Database. The sample coverage for both indices is reported in Table 2.

Daily returns are computed as close-to-close percentage log returns:

$$r_t = 100(\log P_t - \log P_{t-1}), \tag{53}$$

where P_t denotes the daily closing price. The construction of intraday returns and realized variance follows Section 3.1. Specifically, the high-frequency measures use five-minute log returns sampled during regular United States trading hours, from 9:30 a.m. to 4:00 p.m. Eastern Time. The five-minute frequency balances the retention of intraday price variation against market microstructure noise in liquid equity markets. Finer sampling can amplify microstructure noise, whereas coarser sampling reduces the intraday resolution available for identifying jumps; five-minute realized variance is also a widely used benchmark across asset classes [15,16,51]. Using the same sampling grid for both indices and both decompositions keeps the CPR–ABD comparison independent of differences in sampling frequency. The number of available intraday returns on day t is denoted by M_t and is typically 78. Equation (2) is applied with $M = M_t$ to obtain regular-hours realized variance. The CPR- and ABD-based continuous and jump components are constructed from the same intraday observations using the procedures introduced in Section 3.1.

The database supplies five-minute transaction prices directly. The series are complete within each actual trading day, so no interpolation or imputation is applied. The samples follow the trading calendar for each index: weekends and full holidays are excluded, while

early-close sessions retain their shorter schedules and may contain fewer than 78 observations. Overnight price movements are not included in the intraday series.

Figure 2 plots the S&P 500 daily returns and realized variance. The series display the familiar clustering of quiet and turbulent periods, with large realized-variance spikes during the global financial crisis, the coronavirus disease 2019 (COVID-19) market disruption, and later episodes of elevated market uncertainty. The figure shows that large absolute daily returns are generally associated with elevated realized variance, while realized variance also contains substantial intraday variation that is not captured by close-to-close returns alone. This difference motivates the use of high-frequency realized measures in the volatility specification.

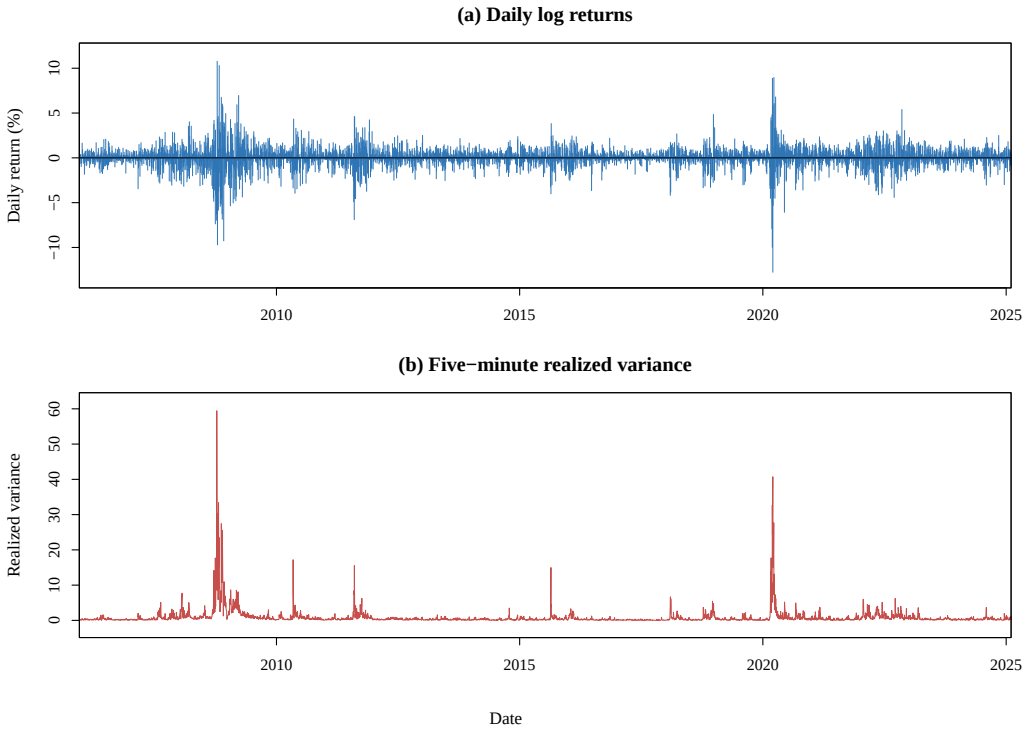


Figure 2. Standard & Poor's 500 (S&P 500) daily returns and realized variance. Notes: Panel (**a**) plots percentage close-to-close log returns. Panel (**b**) plots regular-hours realized variance constructed from five-minute intraday returns as defined in Section 3.1.

Table 2 reports the sample coverage, distributional properties, and jump characteristics of the two indices. Panel A shows that the standard deviation of daily returns is 1.221% for the S&P 500 and 1.360% for the NASDAQ 100. Both return series exhibit substantial excess kurtosis, with values of 12.401 and 7.060, respectively, supporting the use of Student-t innovations. The first-order autocorrelation of log RV_t is 0.829 for the S&P 500 and 0.793 for the NASDAQ 100, indicating pronounced realized-volatility persistence in both samples.

Panel B distinguishes the frequency of detected jump days from the aggregate contribution of jump variation. ABD identifies jump days more frequently for both indices, with frequencies of 42.66–43.35%, compared with 24.30–26.28% under CPR. The corresponding jump shares are 6.08–6.39% under CPR and 8.56–9.24% under ABD. Thus, the difference in detection frequency does not translate proportionally into the aggregate amount of jump variation. The relatively modest aggregate jump shares are consistent with evidence that jumps are episodic and account for a limited fraction of total realized variation [23,24].

For the common S&P 500 likelihood sample of 4266 observations, jumps are identified on 22% of days by both procedures, while 4.0% of days are classified as jump days only by CPR and 21% only by ABD; the remaining 53% are classified as non-jump days by

both procedures. The corresponding Jaccard index and Cohen's κ are both 0.47. The CPR- and ABD-based jump components have a correlation of 0.76, compared with 0.99 for their continuous components. The continuous estimates are therefore closely aligned, whereas the jump classifications retain substantial differences, particularly through the additional jump days identified by ABD.

Table 2. Sample coverage, distributional properties, and jump characteristics.

Panel A: Sample Coverage and Distributional Properties					
Asset	**Sample Period**	**Obs.**	**SD** (r_t)	**ExKurt** (r_t)	**ACF$_1$** $(\log RV_t)$
S&P 500	2005-12-13–2025-02-07	4817	1.221	12.401	0.829
NASDAQ 100	2004-03-05–2025-02-07	5267	1.360	7.060	0.793
Panel B: Jump Incidence and Contribution to Realized Variance					
Asset		**CPR**		**ABD**	
		Jump Days (%)	**Jump Share (%)**	**Jump Days (%)**	**Jump Share (%)**
S&P 500		26.28	6.39	42.66	8.56
NASDAQ 100		24.30	6.08	43.35	9.24

Notes: Daily returns are expressed in percentage points. The abbreviations SD, ExKurt, and ACF$_1$ denote standard deviation, excess kurtosis, and first-order autocorrelation, respectively. CPR and ABD denote the Corsi–Pirino–Renò and Andersen–Bollerslev–Dobrev procedures. Both jump decompositions use the significance level $\alpha_J = 5\%$. Jump days is the percentage of sample days classified as containing jumps by the corresponding procedure. Jump share is defined as $100 \sum_t J_t^{(d)} / \sum_t RV_t$, where $d \in \{\text{CPR}, \text{ABD}\}$. Realized variance and its jump components are constructed from five-minute returns observed during regular trading hours.

The two procedures therefore provide nearly the same continuous signal but materially different jump classifications. Agreement in the persistence of C_t supports the stability of the long-run input across procedures. Because the jump measures differ more substantially, the cross-procedure comparison of the short-run jump coefficient provides a more demanding robustness check, while differences in out-of-sample performance are interpreted separately as decomposition sensitivity.

To evaluate the persistence contrast directly, Table 3 reports component-specific autocorrelations, AR(1) coefficients, and implied half-lives over the common 4266-observation likelihood sample.

Table 3. Persistence diagnostics for continuous and jump variation.

Decomposition	**Series**	**ACF$_1$**	**ACF$_5$**	**ACF$_{22}$**	**AR(1)**	**Half-Life (Days)**
CPR	$\log C_t$	0.840	0.697	0.495	0.840	3.98
CPR	$\log(1 + J_t)$	0.081	0.068	0.014	0.081	0.28
ABD	$\log C_t$	0.847	0.703	0.499	0.847	4.17
ABD	$\log(1 + J_t)$	0.116	0.048	0.072	0.116	0.32

Notes: All statistics use the aligned likelihood sample. The transformation $\log(1 + J_t)$ retains zero-jump days. Half-life is measured in trading days and calculated as $\log(0.5) / \log(\widehat{\rho})$, where $\widehat{\rho}$ is the reported AR(1) slope; it is an AR(1)-implied persistence summary.

The persistence contrast is pronounced under both decompositions. The first-order autocorrelation of $\log C_t$ is approximately 0.84, and its 22-day autocorrelation remains close to 0.50. By comparison, the autocorrelations of $\log(1 + J_t)$ are small at the first lag and remain close to zero at longer lags. The AR(1)-implied half-life is approximately four trading days for continuous variation and less than one trading day for jump variation. These diagnostics provide the empirical basis for associating the more persistent continuous component with the MIDAS term and the less persistent jump component with the short-run channel in this sample.

Taken together, Figure 2 and Tables 2 and 3 establish the empirical features underlying the model specification. Daily returns are heavy-tailed, while log realized variance is

strongly persistent in both indices. Jump variation is also economically meaningful, although its measured incidence and aggregate contribution differ between the CPR and ABD procedures. These features motivate the corresponding elements of the proposed model: Student-t innovations accommodate the non-Gaussian return distribution, the persistent continuous component enters the long-run MIDAS term, and jump variation is assigned to the short-run volatility recursion. The following section examines whether this allocation is supported by the estimated model dynamics and whether it improves volatility and tail-risk forecasts.

5. Results and Discussion

The empirical analysis proceeds in four stages. First, the full-sample estimates are used to examine whether the continuous–jump allocation is supported by the internal dynamics of the model. Second, the OOS analysis evaluates whether this allocation improves volatility forecasts across horizons and volatility states. Third, the value-at-risk (VaR) analysis studies whether any statistical forecasting gains translate into more reliable tail-risk estimates. Finally, the robustness analysis varies the information timing, the forecast evaluation window, and the asset under study to determine the scope of the main findings. Table 4 summarizes the information structure and empirical role of the competing models used throughout these analyses.

Table 4. Models included in the empirical comparison.

Abbreviation	Full Name	Model Feature
GJR-GARCH	Glosten–Jagannathan–Runkle GARCH model	Return-only benchmark with asymmetric responses to negative shocks.
GM-RV	GARCH-MIDAS with realized variance	Undecomposed RV_t drives the long-run MIDAS component.
GM-C(CPR) GM-C(ABD)	GARCH-MIDAS with continuous variation	C_t drives the long-run component, while jump variation is excluded.
GM-CJ(CPR) GM-CJ(ABD)	GARCH-MIDAS with continuous and jump variation	C_t enters the long-run component and demeaned J_t enters the short-run recursion.
HAR-CJ(CPR) HAR-CJ(ABD)	HAR model with continuous and jump variations	Uses daily, weekly, and monthly continuous and jump measures.
RGARCH	Realized GARCH model	Links RV_t to conditional variance through a measurement equation.

Notes: GARCH denotes generalized autoregressive conditional heteroskedasticity, MIDAS denotes mixed data sampling, and HAR denotes heterogeneous autoregressive. CPR and ABD denote the Corsi–Pirino–Renò and Andersen–Bollerslev–Dobrev jump-detection procedures used to construct the continuous component C_t and the jump component J_t. Specifications reported with both suffixes contribute two variants, giving nine models in the full OOS comparison.

The model set is designed to distinguish the value of the complete continuous–jump allocation from the separate contributions of its two components. Within the GARCH-MIDAS family, comparing GM-RV with GM-C shows whether removing detected jump variation from the long-run input preserves the persistent information contained in aggregate realized variance. Comparing GM-C with GM-CJ then isolates the incremental contribution of returning jump variation through the short-run recursion, while the direct GM-RV–GM-CJ comparison evaluates the complete reallocation against an undecomposed realized-variance specification. The CPR and ABD variants indicate whether these conclusions depend on the jump-detection procedure. GJR-GARCH, HAR-CJ, and Realized GARCH provide complementary return-only, decomposition-based, and measurement-

equation benchmarks, respectively. Accordingly, evidence for the proposed mechanism is assessed through both the nested GM-C–GM-CJ comparison and the broader forecasting performance of GM-CJ relative to the full benchmark set.

5.1. Full-Sample Estimation and Tests of the Jump Component

Table 5 reports the full-sample estimates for the S&P 500 under the $t-2$ information convention. (For parameters estimated in the interior of the admissible parameter space, parentheses contain Bollerslev–Wooldridge-type QML-robust sandwich standard errors computed from the standardized Student-t quasi-likelihood [48]. The estimates of α, marked by †, lie at the admissible lower boundary. Francq and Zakoïan [52] develop a one-sided expansion for boundary GARCH quasi-maximum-likelihood estimators (QMLEs) based on right derivatives and show that the limiting estimator is a projection of a Gaussian vector onto a convex cone. Motivated by this construction, the parenthesized quantities beneath α are robust sandwich scales computed using feasible right-hand numerical derivatives and a centered score covariance matrix at the constrained optimum. They measure the local scale of the Gaussian component entering the cone-projection limit, rather than conventional standard errors of the constrained estimator, and are not used for Wald inference, symmetric confidence intervals, p-values, or significance classification [50].) The GARCH-MIDAS specifications use $N = 22$ and $K = 24$ non-overlapping blocks. Their effective estimation sample contains 4266 observations, compared with 4795 observations for GJR-GARCH; raw likelihood values should therefore be compared only among models estimated over the same sample.

The symmetric short-run news coefficient α is estimated at its admissible lower boundary of zero in every specification. This is a valid constrained QML outcome. Francq and Zakoïan [52] establish QML asymptotics for GARCH processes when one or more coefficients are zero and show that the limiting distribution is generally a projection of a Gaussian vector onto a convex cone. The consequence is nonstandard inference at the boundary, not an invalid fitted volatility process; conventional two-sided Wald inference for α is therefore inappropriate [50]. Zero estimates of the symmetric news coefficient have also been reported in daily-frequency GARCH-MIDAS applications by Conrad and Kleen [20]. In the present application, the same boundary solution arises in the return-only GJR-GARCH model and in all five GARCH-MIDAS specifications, while every fitted conditional-variance path remains strictly positive. The result is therefore neither specific to the continuous–jump decomposition nor caused by the MIDAS timing convention. The estimates and robust uncertainty measures for parameters in the interior remain informative, although joint inference involving a boundary parameter requires the nonstandard methods used below.

Economically, $\alpha = 0$ implies that the data do not support an additional symmetric contemporaneous response to squared return shocks once the asymmetric response, volatility persistence, the long-run component, and, where applicable, the jump channel are included. Negative return shocks have an immediate coefficient $\alpha + \gamma = \gamma$, and γ is positive and precisely estimated in every specification, while β preserves substantial short-run persistence. Persistence is 0.982 in GJR-GARCH and ranges from 0.890 to 0.894 in the GARCH-MIDAS specifications. The estimated degrees of freedom, between 5.79 and 6.62, also indicate that the standardized Student-t distribution continues to play an important role in accommodating large return innovations. Re-estimating all six specifications with α fixed at zero reproduces the boundary solutions: the largest absolute difference in log-likelihood is 1.0×10^{-7}, and the largest absolute difference between the fitted conditional-variance paths is 0.00184.

Table 5. Full-sample parameter estimates for the S&P 500.

	GJR-GARCH	GM-RV	GM-C (CPR)	GM-CJ (CPR)	GM-C (ABD)	GM-CJ (ABD)
Panel A: Return and short-run components						
μ	0.064 ***	0.064 ***	0.064 ***	0.064 ***	0.064 ***	0.064 ***
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
ω	0.021 ***	–	–	–	–	–
	(0.004)					
$\alpha^{\dagger}$	0.000	0.000	0.000	0.000	0.000	0.000
	(0.011)	(0.035)	(0.034)	(0.041)	(0.034)	(0.042)
γ	0.234 ***	0.284 ***	0.284 ***	0.270 ***	0.284 ***	0.274 ***
	(0.027)	(0.028)	(0.028)	(0.028)	(0.028)	(0.028)
β	0.865 ***	0.748 ***	0.752 ***	0.755 ***	0.751 ***	0.753 ***
	(0.013)	(0.023)	(0.023)	(0.021)	(0.023)	(0.021)
Panel B: Long-run and jump components						
m	–	0.808 ***	0.872 ***	0.811 ***	0.894 ***	0.828 ***
		(0.095)	(0.100)	(0.095)	(0.100)	(0.098)
λ	–	0.894 ***	0.884 ***	0.781 ***	0.874 ***	0.765 ***
		(0.057)	(0.057)	(0.072)	(0.056)	(0.075)
w_2	–	25.948 ***	23.643 ***	24.099 ***	24.082 ***	22.426 ***
		(9.539)	(8.729)	(7.650)	(8.835)	(7.338)
ζ	–	–	–	0.381 ***	–	0.325 ***
				(0.143)		(0.133)
Panel C: Distribution, variance allocation, and fit						
ν	5.792 ***	6.480 ***	6.477 ***	6.617 ***	6.473 ***	6.588 ***
	(0.484)	(0.650)	(0.649)	(0.683)	(0.648)	(0.677)
Persistence	0.982	0.890	0.894	0.890	0.893	0.890
VR (%)	–	56.93	55.92	43.35	56.15	41.83
LogLik	−6227.950	−5577.810	−5578.270	−5573.570	−5578.150	−5573.700
N_{eff}	4795	4266	4266	4266	4266	4266

Notes: Persistence is $\alpha + \gamma/2 + \beta$. N_{eff} is the effective estimation sample. A dagger indicates that α is estimated at its admissible lower boundary of zero. Asterisks *** denote significance at the 1% level. The significance markers for ζ are based on the conditional parametric-bootstrap likelihood-ratio tests reported in Table 6.

Panel B first examines whether the continuous component preserves the long-run information contained in aggregate realized variance. The estimates of m, λ, and w_2 are statistically significant in every GARCH-MIDAS specification. Replacing $\log RV_t$ with $\log C_t$ changes neither the long-run parameter profile nor the maximized likelihood materially: relative to GM-RV, the log-likelihood decreases by only 0.46 for GM-C(CPR) and 0.34 for GM-C(ABD). Following Conrad and Kleen [20], panel C complements these estimates with the variance ratio $\text{VR} = 100\ \text{Var}(\log \tau_t)/\ \text{Var}(\log \sigma_t^2)$, where $\sigma_t^2 = \tau_t g_t$. The ratio is computed at the daily frequency over the common GARCH-MIDAS estimation sample and gauges the relative importance of the long-run component. A higher VR indicates that a larger proportion of overall conditional-variance movements can be traced to τ_t. The VR results reinforce the likelihood comparison: GM-RV records a ratio of 56.93%, compared with 55.92% and 56.15% for the CPR- and ABD-based continuous-only specifications. Thus, removing detected jump variation from the MIDAS input preserves almost all of the persistent variation attributed to aggregate realized variance. When jump variation is reintroduced through the short-run recursion, the ratio declines to 43.35% under CPR and 41.83% under ABD. This reduction is consistent with the complete CJ specification assigning a larger part of conditional-variance fluctuations to the short-run component rather than forcing them into τ_t. The variance-ratio shift describes the allocation of conditional-variance movements within the fitted model, while the persistence distinction itself is established directly by the component diagnostics in Table 3. Figure 3 further examines whether the temporal weighting of the long-run information changes across these specifications.

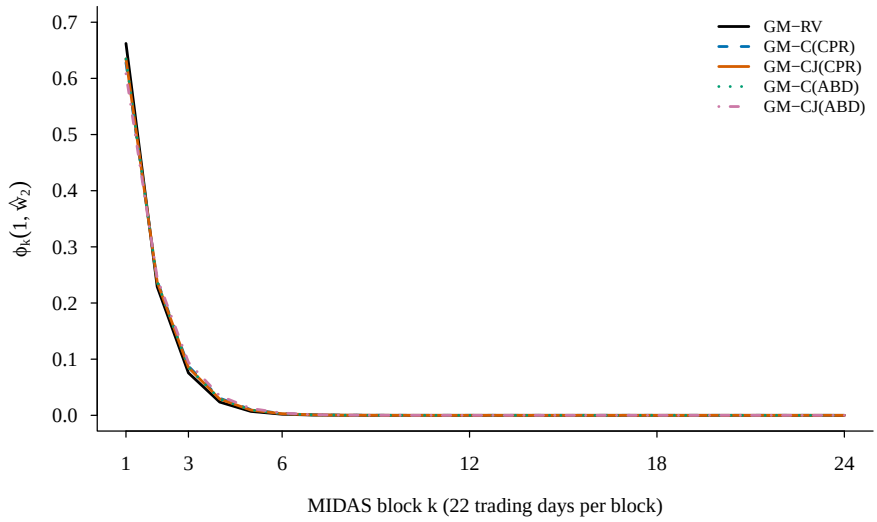


Figure 3. Estimated mixed-data-sampling (MIDAS) weighting functions. Notes: The figure plots the restricted Beta weights $\phi_k(1, \widehat{w}_2)$ implied by the full-sample estimates. Each MIDAS block contains $N = 22$ non-overlapping daily observations, and the maximum lag length is $K = 24$ blocks. Block $k = 1$ is the most recent block available under the $t - 2$ information convention.

The estimated weighting functions are monotonically and rapidly declining in every specification. The most recent 22-day block receives between 60.8% and 66.2% of the total MIDAS weight, while the first three blocks together receive between 94.7% and 96.7%. More than 99.8% of the total weight is assigned to the first six blocks, equivalent to 132 trading days. The nominal 528-day window therefore defines the maximum support, whereas the effective memory is concentrated within the most recent 132 trading days; more distant blocks receive negligible weight. The five curves are also remarkably similar. Decomposing realized variance and introducing the jump channel therefore affect the loading and the allocation of conditional-variance movements more strongly than they affect the estimated decay pattern of the long-run information.

Comparisons between GM-CJ and the corresponding GM-C specifications isolate the effect of introducing jump variation into the short-run recursion. The maximized log-likelihood increases by 4.70 under CPR and by 4.45 under ABD. Meanwhile, the estimated long-run loading $\widehat{\lambda}$ declines from 0.8841 to 0.7810 for CPR and from 0.8744 to 0.7646 for ABD. This joint pattern indicates that allowing jumps to affect the short-run dynamics improves the in-sample fit while reducing the loading placed on the persistent continuous component, consistent with the intended separation of transitory jump variation from longer-run volatility movements.

The fitted GM-CJ recursions also satisfy the sample sufficient condition for positivity. The margins $\widehat{c} - \widehat{\zeta}\overline{J}$ are 0.089 under CPR and 0.086 under ABD, while the corresponding full-sample minima of $\widehat{g}_t$ are 0.38 and 0.37. Across the 778 rolling re-estimations for each decomposition, the minimum fitted $\widehat{g}_t$ is 0.31 under CPR and 0.32 under ABD. Thus, both the analytical margin and the fitted short-run paths remain strictly positive under the two decompositions.

As set out in Section 3, the nested comparison between GM-C and GM-CJ tests $H_0 : \zeta = 0$ against the one-sided alternative $H_1 : \zeta > 0$. Table 6 reports the one-sided Wald p-value p_W, the likelihood-ratio statistic LR_ζ, its boundary-mixture p-value p_{mix}, and the conditional parametric-bootstrap p-value p_{boot}. (Here $\widehat{\zeta}$ is the constrained maximum-likelihood estimate of the nonnegative jump loading in Equation (31); p_W, LR_ζ, p_{mix}, and p_{boot} are defined in Equations (47), (48), (50), and (52), respectively. The mixture calculation uses the $\frac{1}{2}\chi^2_0 + \frac{1}{2}\chi^2_1$ reference distribution. For the bootstrap, both the restricted and unrestricted models are re-estimated under their original parameter constraints in every

replication. The requested/successful/unsuccessful replication counts are 499/499/0 for CPR and 999/999/0 for ABD.) Because $\zeta = 0$ is the boundary null and α is also estimated on an active boundary in the present application, inference is based primarily on p_{boot}. The Wald and boundary-mixture results are reported as complementary evidence.

Table 6. Tests of the jump-component coefficient for the S&P 500 index.

Decomposition	$\widehat{\zeta}$	p_W	LR_ζ	p_{mix}	p_{boot}
CPR	0.3808	0.0040	9.4078	0.0011	0.0020
ABD	0.3252	0.0073	8.8930	0.0014	0.0030

Notes: CPR and ABD denote the Corsi–Pirino–Renò and Andersen–Bollerslev–Dobrev decompositions, respectively.

The conditional bootstrap, which provides the primary inference in the presence of the active boundary parameters, rejects $H_0 : \zeta = 0$ under both decompositions: $p_{\text{boot}} = 0.0020$ for CPR and 0.0030 for ABD. The supplementary procedures yield the same conclusion. The one-sided Wald p-values are 0.0040 and 0.0073, while the boundary-mixture likelihood-ratio p-values are 0.0011 and 0.0014, respectively. Evidence for a positive short-run jump loading is therefore present under both decompositions and does not depend on a single inferential procedure.

The coefficient estimates also give the magnitude of the short-run jump response. Multiplying $\widehat{\zeta}$ by the sample standard deviation of J_t gives increases of 0.15 under CPR and 0.14 under ABD in the short-run multiplier g_t. Evaluated at the normalization point $g_t = 1$, these correspond to increases of approximately 15% and 14%, respectively. Since conditional volatility is $\sqrt{\tau_t g_t}$, the associated increases in conditional volatility are approximately 7.1% and 6.9%, holding τ_t and the remaining short-run inputs fixed.

Taken together with the near preservation of the long-run structure when GM-RV is replaced by GM-C, these results support the complete CJ allocation: continuous variation retains the persistent information assigned to the long-run component, while jump variation contributes separately through the short-run recursion. This remains an in-sample statement about the internal structure of the model rather than a claim of general predictive dominance. The broader comparison with GJR-GARCH, GM-RV, HAR-CJ, and Realized GARCH is therefore conducted through OOS forecasts in the next subsection.

5.2. Volatility Forecasting Evaluation

We evaluate the competing models using fixed-length rolling windows. The OOS period starts in January 2022, and forecasts are generated for $h = 1, 5, 10$, and 22 trading days. All models are re-estimated at each forecast origin using only information available by the forecast date. For the GARCH-MIDAS specifications, the baseline $t - 2$ convention defined in Section 3 is adopted.

The evaluation window covers the post-pandemic inflation episode and the Federal Reserve's rapid monetary tightening. It also contains several episodes of elevated market volatility, providing a demanding setting for assessing the forecasting content of jump variation. The alternative-window results in Section 5.4 show how the magnitude of the forecast gains changes with the evaluation period.

Since actual volatility is unobservable, we adopt the scaled realized variance measure of Hansen and Lunde [53] as its proxy. The scaled realized variance is defined as

$$\widetilde{RV}_t := \widehat{\eta}_{\text{RV}} RV_t, \qquad \widehat{\eta}_{\text{RV}} := \frac{\sum_{t=1}^{T} (r_t - \bar{r})^2}{\sum_{t=1}^{T} RV_t}, \qquad t = 1, 2, \ldots, T, \tag{54}$$

where $\bar{r} = T^{-1}\sum_{t=1}^{T} r_t$ is the sample mean of daily returns. This adjustment accounts for price variation during non-trading hours.

Two multi-step forecast targets are considered. The direct point target evaluates conditional variance on day $t+h$, whereas the average cumulative target evaluates mean variance over days $t+1$ through $t+h$. Let $\widehat{\sigma}^2_{t+j|t}$ denote the model-based conditional-variance forecast for day $t+j$ constructed using information available until day t. The realized targets and their corresponding forecasts are

$$\widetilde{V}^D_{t,h} = \widetilde{RV}_{t+h}, \qquad \widehat{V}^D_{t,h} = \widehat{\sigma}^2_{t+h|t}, \tag{55}$$

and

$$\widetilde{V}^A_{t,h} = \frac{1}{h}\sum_{j=1}^{h} \widetilde{RV}_{t+j}, \qquad \widehat{V}^A_{t,h} = \frac{1}{h}\sum_{j=1}^{h} \widehat{\sigma}^2_{t+j|t}, \tag{56}$$

where the superscripts D and A denote the direct point and average cumulative targets, respectively.

The main forecasting comparison uses two representative loss functions: the mean squared error (MSE) and the quasi-likelihood (QLIKE). For $s \in \{D, A\}$, forecast accuracy is measured by

$$\mathrm{MSE}^s_h = \frac{1}{T_h}\sum_{t=1}^{T_h}\left(\widetilde{V}^s_{t,h} - \widehat{V}^s_{t,h}\right)^2, \tag{57}$$

$$\mathrm{QLIKE}^s_h = \frac{1}{T_h}\sum_{t=1}^{T_h}\left[\log \widehat{V}^s_{t,h} + \frac{\widetilde{V}^s_{t,h}}{\widehat{V}^s_{t,h}}\right], \tag{58}$$

where T_h is the number of OOS forecasts at horizon h. The reported QLIKE values are normalized by subtracting $T_h^{-1}\sum_{t=1}^{T_h}[\log \widetilde{V}^s_{t,h} + 1]$. This term is common to all models in a given evaluation sample, so the normalization leaves model rankings and loss differentials unchanged. Both MSE and QLIKE are robust to noise in a conditionally unbiased volatility proxy. Within this class, MSE is unique in depending solely on the forecast error, whereas QLIKE is unique in depending solely on the standardized forecast error [54]. MSE penalizes over- and underprediction symmetrically but is relatively sensitive to extreme observations. QLIKE is asymmetric, assigns a larger penalty to volatility underprediction, and is less sensitive to extreme observations.

As additional checks, we consider the mean absolute error (MAE) and the robust loss function (RLF), taken as the $\beta = -1$ member of the loss family in Patton [54]:

$$\mathrm{MAE}^s_h = \frac{1}{T_h}\sum_{t=1}^{T_h}\left|\widetilde{V}^s_{t,h} - \widehat{V}^s_{t,h}\right|, \tag{59}$$

$$\mathrm{RLF}^s_h = \frac{1}{T_h}\sum_{t=1}^{T_h}\left[\widehat{V}^s_{t,h} - \widetilde{V}^s_{t,h} + \widetilde{V}^s_{t,h}\log\left(\frac{\widetilde{V}^s_{t,h}}{\widehat{V}^s_{t,h}}\right)\right]. \tag{60}$$

MAE provides a complementary absolute-error criterion, consistent with the use of multiple evaluation metrics in recent forecasting work [55], while RLF supplies an additional proxy-robust criterion with a different curvature from MSE and QLIKE. The corresponding results are reported in Appendix D.

Statistical differences in predictive ability are assessed using the 90% model confidence set (MCS) of Hansen et al. [56]. For each combination of forecast target, horizon, and loss function, this procedure begins with the complete set of competing models and iteratively removes the worst-performing model until the null hypothesis of equal predictive ability can no longer be rejected for the remaining models, which constitute the MCS. We compute

the MCS using 5000 stationary-bootstrap replications to accommodate serial dependence in the loss differentials. A shaded entry in the tables identifies a model that remains in the 90% MCS. Point losses and loss differentials describe the numerical magnitude of forecast differences, whereas MCS membership indicates statistical competitiveness within a given target, horizon, and loss function. For each decomposition and forecast setting, we also apply the same MCS procedure to the nested GM-CJ and GM-C pair. These comparisons address distinct questions: the jump-coefficient tests concern the short-run jump effect, log-likelihood differences concern in-sample fit, mean loss differences quantify numerical out-of-sample accuracy, and the focused MCS provides a formal comparison of relative out-of-sample predictive ability.

Table 7 presents the broad OOS comparison. Panel A reports direct point forecasts and panel B reports average cumulative forecasts, with all four horizons retained to show how relative performance changes with h.

Table 7. Out-of-sample volatility forecast losses for the S&P 500.

Model	MSE				QLIKE			
	$h = 1$	$h = 5$	$h = 10$	$h = 22$	$h = 1$	$h = 5$	$h = 10$	$h = 22$
Panel A: Direct Point Forecasts								
HAR-CJ(CPR)	1.460	1.787	1.914	1.648	0.409	0.491	0.545	0.476
HAR-CJ(ABD)	1.439	1.750	1.887	1.641	0.407	0.479	0.534	0.477
GJR-GARCH	1.204	1.703	1.956	1.857	0.203	0.267	0.328	0.381
GM-RV	1.191	1.477	1.633	1.499	0.187	0.257	0.313	0.324
GM-C(CPR)	1.203	1.494	1.646	1.505	0.187	0.259	0.314	0.324
GM-CJ(CPR)	1.183	**1.455**	**1.592**	1.446	0.188	**0.256**	**0.309**	0.319
GM-C(ABD)	1.193	1.480	1.634	1.488	**0.187**	0.258	0.315	0.324
GM-CJ(ABD)	1.255	1.493	1.609	**1.431**	0.190	0.259	0.312	**0.319**
Realized-GARCH	**1.140**	1.561	1.728	1.554	0.302	0.363	0.394	0.384
Panel B: Average Cumulative Forecasts								
HAR-CJ(CPR)	1.460	0.996	0.897	0.747	0.409	0.346	0.340	0.305
HAR-CJ(ABD)	1.439	0.967	0.862	0.725	0.407	0.338	0.329	0.298
GJR-GARCH	1.204	0.812	0.798	0.798	0.203	0.134	0.132	0.144
GM-RV	1.191	0.689	0.585	0.497	0.187	0.121	0.117	0.117
GM-C(CPR)	1.203	0.704	0.601	0.510	0.187	0.122	0.118	0.118
GM-CJ(CPR)	1.183	**0.671**	**0.558**	**0.457**	0.188	**0.121**	**0.116**	**0.113**
GM-C(ABD)	1.193	0.692	0.589	0.497	**0.187**	0.121	0.118	0.118
GM-CJ(ABD)	1.255	0.728	0.600	0.475	0.190	0.123	0.118	0.115
Realized-GARCH	**1.140**	0.731	0.662	0.571	0.302	0.237	0.216	0.184

Notes: Lower values are better. Bold values are the lowest loss in each horizon–loss column. Shaded cells denote membership in the 90% model confidence set (MCS), constructed using 5000 stationary-bootstrap replications.

For direct point forecasts, GM-CJ(CPR) attains the lowest MSE and QLIKE at $h = 5$ and $h = 10$ and belongs to the corresponding 90% MCSs. At $h = 22$, GM-CJ(ABD) records the lowest loss under both criteria, while GM-CJ(CPR) remains in the QLIKE MCS. At $h = 1$, other realized-measure models attain the lowest losses, although GM-CJ(CPR) remains in the 90% MCS under both criteria. The lowest direct point losses for the GM-CJ specifications are therefore concentrated at the multi-day horizons.

For average cumulative forecasts, GM-CJ(CPR) achieves the lowest MSE and QLIKE at $h = 5, 10$, and 22 and is retained in the 90% MCS under both loss functions at all three horizons. Within the GARCH-MIDAS family, adding jump variation to the short-run recursion also lowers the multi-day losses relative to GM-C(CPR). Thus, the lowest multi-day losses for GM-CJ(CPR) occur for the average cumulative target.

The focused out-of-sample comparison provides more consistent evidence of incremental predictive improvement under CPR. GM-CJ is retained in all 16 focused MCSs, compared with five for GM-C, and records the lower mean loss in 14 comparisons. Un-

der ABD, each model is retained in 13 focused MCSs, with GM-CJ recording the lower mean loss in six comparisons. Thus, evidence for a positive short-run jump loading is present under both decompositions, whereas the incremental out-of-sample evidence is stronger under CPR and depends on the forecast setting.

To examine performance across market conditions, we group forecasts into low-, medium-, and high-volatility states following Conrad and Kleen [20]. The states are defined by whether the corresponding scaled realized-variance target is at or below the 25th percentile, between the 25th and 75th percentiles, or above the 75th percentile, respectively (denoted by 25/75). Table 8 reports the resulting MSE and QLIKE losses in panels A and B; each panel retains the direct point and average cumulative targets at all four horizons.

The state-based results clarify where the loss differences arise. In the low-volatility state, the HAR-CJ benchmarks record the smallest losses. In the medium-volatility state, GM-CJ(CPR) attains the lowest QLIKE for average cumulative forecasts at $h = 5$, remains close to the minimum at $h = 10$, and is retained in the 90% MCS at every horizon. GM-CJ(ABD) records the lowest QLIKE for average cumulative forecasts at $h = 10$ and $h = 22$.

In the high-volatility state, GM-CJ(CPR) attains the lowest MSE for average cumulative forecasts at $h = 5, 10$, and 22. The GM-CJ specifications also record the lowest QLIKE for direct point forecasts at the longer horizons: GM-CJ(CPR) at $h = 10$ and GM-CJ(ABD) at $h = 22$. GM-CJ(CPR) records the lowest QLIKE for average cumulative forecasts at $h = 22$ and remains close to the minimum at $h = 10$. It is retained in the 90% MCS in all 16 high-volatility comparisons. The point-loss minima are concentrated at multi-day horizons, while the MCS results show statistical competitiveness across the full set of high-volatility comparisons. With either 20/80 or 30/70 cutoffs, GM-CJ(CPR) continues to record lower average-cumulative MSE than GM-C(CPR) at $h = 5, 10$, and 22; the remaining patterns vary across forecast targets, loss functions, and decompositions. Appendix C reports the corresponding cutoff comparison.

5.3. Value-At-Risk Forecasting

Beyond point forecast accuracy, a volatility model used for risk management must also provide reliable forecasts of downside risk. We therefore evaluate one-day-ahead value-at-risk (VaR), where the VaR at tail probability α is the conditional lower-tail return threshold below which returns should fall with probability α. At each OOS forecast origin, the parametric VaR forecast is

$$\widehat{\mathrm{VaR}}_{t+1|t}(\alpha) = \widehat{\mu}_{t+1|t} + \widehat{\sigma}_{t+1|t} q_{\alpha}^{std}(\widehat{\nu}_t), \tag{61}$$

where $\widehat{\mu}_{t+1|t}$ and $\widehat{\sigma}^2_{t+1|t}$ are the one-day-ahead conditional mean and variance forecasts, respectively, and $q_{\alpha}^{std}(\widehat{\nu}_t)$ is the lower-tail quantile of the variance-standardized Student-t distribution used in estimation:

$$q_{\alpha}^{std}(\nu) = \sqrt{\frac{\nu - 2}{\nu}}\, F_{\nu}^{-1}(\alpha), \tag{62}$$

where $F_{\nu}^{-1}(\alpha)$ is the α-quantile of a Student-t distribution with ν degrees of freedom. The factor $\sqrt{(\nu-2)/\nu}$ rescales that quantile so that the innovation has unit variance. This construction makes the VaR forecast consistent with both the conditional-variance dynamics and the heavy-tailed innovation distribution used in estimation.

Table 8. Out-of-sample volatility forecast losses across volatility states for the S&P 500.

Panel A: MSE												
	Low-Volatility State				**Medium-Volatility State**				**High-Volatility State**			
Model	$h = 1$	$h = 5$	$h = 10$	$h = 22$	$h = 1$	$h = 5$	$h = 10$	$h = 22$	$h = 1$	$h = 5$	$h = 10$	$h = 22$
Direct Point Forecasts												
HAR-CJ(CPR)	0.014	**0.050**	**0.076**	**0.106**	**0.178**	**0.182**	**0.188**	**0.164**	5.454	6.719	7.215	6.165
HAR-CJ(ABD)	**0.014**	0.055	0.082	0.109	0.182	0.187	0.205	0.168	5.367	6.554	7.065	6.127
GJR-GARCH	0.139	0.393	0.455	0.751	0.420	0.731	0.665	0.846	3.831	4.946	6.046	4.990
GM-RV	0.114	0.307	0.343	0.409	0.414	0.652	0.665	0.853	**3.815**	**4.289**	4.868	3.888
GM-C(CPR)	0.117	0.315	0.354	0.426	0.423	0.671	0.684	0.882	3.841	4.312	4.868	**3.834**
GM-CJ(CPR)	0.112	0.311	0.362	0.443	0.390	0.595	0.589	0.734	3.831	4.309	**4.836**	3.879
GM-C(ABD)	0.114	0.310	0.345	0.412	0.413	0.653	0.657	0.835	3.821	4.296	4.885	3.874
GM-CJ(ABD)	0.113	0.314	0.365	0.443	0.440	0.612	0.600	0.699	4.017	4.422	4.878	3.888
Realized-GARCH	0.028	0.155	0.250	0.427	0.202	0.270	0.319	0.171	4.120	5.535	6.031	5.454
Average Cumulative Forecasts												
HAR-CJ(CPR)	0.014	**0.019**	**0.018**	**0.016**	**0.178**	0.178	0.172	0.182	5.454	3.602	3.230	2.612
HAR-CJ(ABD)	**0.014**	0.019	0.019	0.017	0.182	0.176	0.170	0.183	5.367	3.491	3.092	2.518
GJR-GARCH	0.139	0.125	0.097	0.143	0.420	0.246	0.312	0.356	3.831	2.625	2.475	2.338
GM-RV	0.114	0.096	0.090	0.105	0.414	0.240	0.284	0.342	**3.815**	2.175	1.685	1.201
GM-C(CPR)	0.117	0.099	0.094	0.111	0.423	0.245	0.291	0.352	3.841	2.225	1.733	1.226
GM-CJ(CPR)	0.112	0.102	0.101	0.120	0.390	0.226	0.256	0.295	3.831	**2.124**	**1.623**	**1.118**
GM-C(ABD)	0.114	0.096	0.091	0.107	0.413	0.240	0.283	0.334	3.821	2.186	1.700	1.214
GM-CJ(ABD)	0.113	0.105	0.101	0.122	0.440	0.235	0.262	0.300	4.017	2.333	1.780	1.181
Realized-GARCH	0.028	0.042	0.043	0.054	0.202	**0.157**	**0.169**	**0.181**	4.120	2.562	2.267	1.872
Panel B: QLIKE												
	Low-Volatility State				**Medium-Volatility State**				**High-Volatility State**			
Model	$h = 1$	$h = 5$	$h = 10$	$h = 22$	$h = 1$	$h = 5$	$h = 10$	$h = 22$	$h = 1$	$h = 5$	$h = 10$	$h = 22$
Direct Point Forecasts												
HAR-CJ(CPR)	0.105	**0.151**	**0.193**	**0.255**	0.337	0.281	0.242	0.174	0.856	1.251	1.504	1.302
HAR-CJ(ABD)	**0.101**	0.154	0.197	0.258	0.346	0.285	0.247	0.177	0.835	1.190	1.448	1.298
GJR-GARCH	0.226	0.369	0.457	0.653	0.162	0.159	**0.152**	0.155	0.262	0.379	0.550	0.561
GM-RV	0.216	0.381	0.441	0.501	0.145	0.148	0.170	0.190	0.242	0.351	0.472	0.415
GM-C(CPR)	0.221	0.388	0.451	0.512	**0.143**	0.149	0.171	0.192	**0.240**	**0.348**	0.464	0.399
GM-CJ(CPR)	0.213	0.395	0.470	0.540	0.144	0.138	0.157	0.175	0.253	0.352	**0.452**	0.386
GM-C(ABD)	0.218	0.384	0.446	0.507	0.144	0.148	0.169	0.188	0.240	0.350	0.474	0.411
GM-CJ(ABD)	0.213	0.404	0.480	0.545	0.148	**0.138**	0.156	0.172	0.250	0.356	0.457	**0.385**
Realized-GARCH	0.133	0.243	0.359	0.537	0.261	0.211	0.159	**0.096**	0.554	0.783	0.899	0.808
Average Cumulative Forecasts												
HAR-CJ(CPR)	0.105	**0.092**	**0.070**	**0.043**	0.337	0.296	0.280	0.246	0.856	0.700	0.729	0.684
HAR-CJ(ABD)	**0.101**	0.097	0.073	0.046	0.346	0.293	0.278	0.247	0.835	0.669	0.687	0.654
GJR-GARCH	0.226	0.158	0.128	0.123	0.162	0.102	0.109	0.121	0.262	0.172	0.180	0.212
GM-RV	0.216	0.147	0.124	0.110	0.145	0.092	0.096	0.112	0.242	0.153	**0.153**	0.134
GM-C(CPR)	0.221	0.151	0.128	0.116	**0.143**	0.092	0.096	0.111	**0.240**	0.153	0.153	0.134
GM-CJ(CPR)	0.213	0.151	0.134	0.126	0.144	**0.087**	0.087	0.099	0.253	0.159	0.154	**0.130**
GM-C(ABD)	0.218	0.148	0.126	0.113	0.144	0.092	0.096	0.110	0.240	**0.152**	0.155	0.140
GM-CJ(ABD)	0.213	0.156	0.138	0.130	0.148	0.088	**0.087**	**0.098**	0.250	0.161	0.161	0.136
Realized-GARCH	0.133	0.124	0.087	0.071	0.261	0.210	0.195	0.140	0.554	0.406	0.387	0.384

Notes: Lower values are better. Bold values are the lowest loss within each panel, forecast target, state, and horizon. Shaded cells denote membership in the 90% MCS, constructed using 5000 stationary-bootstrap replications. For each forecast target and horizon, low, medium, and high volatility correspond to observations at or below the 25th percentile, between the 25th and 75th percentiles, and above the 75th percentile of the corresponding out-of-sample scaled realized-variance target, respectively.

We evaluate VaR forecasts along two dimensions: predictive accuracy and statistical validity. Predictive accuracy is measured using the quantile loss of González-Rivera et al. [57]:

$$L_{t+1}^{\alpha} = \left(\alpha - I_{t+1}^{\alpha}\right)\left[r_{t+1} - \widehat{\mathrm{VaR}}_{t+1|t}(\alpha)\right], \qquad I_{t+1}^{\alpha} = \mathbf{1}\left\{r_{t+1} < \widehat{\mathrm{VaR}}_{t+1|t}(\alpha)\right\}. \tag{63}$$

where I_{t+1}^{α} is the violation indicator. This loss function asymmetrically penalizes errors by assigning a higher weight of $1-\alpha$ to costly instances where the loss exceeds the predicted VaR. Thus, smaller values indicate more accurate forecasts. As in Section 5.2, we use the 90% MCS to compare the quantile-loss performance of the competing VaR forecasts, as in previous studies [58,59]. This loss-based comparison is reported separately from the coverage diagnostics below.

Statistical validity requires both correct coverage and independently distributed violations. We therefore report the violation rate

$$\widehat{p}_{\alpha} := \frac{1}{T_1}\sum_{t=1}^{T_1} I_{t+1}^{\alpha}, \tag{64}$$

where $T_1 = 778$ is the number of one-day-ahead OOS forecasts. The nominal expected number of violations is $T_1\alpha$, giving 38.90, 7.78, and 3.89 violations at the 5%, 1%, and 0.5% levels, respectively. A valid model should produce a rate close to α. A rate above α indicates too many violations and overestimation of the lower-tail return quantile, whereas a rate below α indicates too few violations and insufficient model forecasts; either departure indicates coverage miscalibration.

The unconditional-coverage (UC) test of Kupiec [60] examines the null hypothesis $\Pr(I_{t+1}^{\alpha} = 1) = \alpha$; its likelihood-ratio statistic LR_{UC} is asymptotically distributed as χ_1^2, a chi-squared distribution with one degree of freedom. The independence (IND) test of Christoffersen [61] assesses whether successive violations are independent. Under the null hypothesis of a correctly specified model, its likelihood-ratio statistic LR_{IND} is also asymptotically χ_1^2 distributed. Finally, the conditional-coverage (CC) test of Christoffersen et al. [62] jointly assesses correct coverage and independence, with the test statistic $LR_{\mathrm{CC}} = LR_{\mathrm{UC}} + LR_{\mathrm{IND}}$ asymptotically following a chi-squared distribution with two degrees of freedom. Accordingly, UC evaluates the violation frequency, IND evaluates violation clustering, and CC evaluates both properties jointly.

Table 9 reports the quantile loss, observed violation count and rate, and UC, CC, and IND p-values for one-day-ahead VaR forecasts at $\alpha = 5\%, 1\%$, and 0.5% levels. Reporting all three tail probabilities allows predictive accuracy to be considered jointly with increasingly stringent tail calibration.

At the 5% tail probability, GM-CJ(ABD) attains the lowest quantile loss, while the loss of GM-CJ(CPR) is nearly identical; both remain in the 90% MCS. GM-CJ(ABD) and GM-CJ(CPR) produce 55 and 56 violations, respectively, compared with a nominal expected count of 38.90. The corresponding rates of 7.07% and 7.20% indicate overestimation of the lower-tail return quantiles at the 5% level, so realized returns fall below the predicted thresholds more often than intended. If used without additional calibration, these forecasts could lead to risk limits or capital buffers that are too low. The UC and CC tests reject correct 5% coverage for the GARCH-MIDAS specifications, whereas the IND test for GM-CJ(CPR) does not reject violation independence. The favorable quantile-loss ranking therefore indicates relative predictive accuracy rather than correct 5% calibration; for GM-CJ(CPR), the departure reflects excessive violation frequency rather than violation clustering.

Table 9. One-day-ahead value-at-risk (VaR) forecasting evaluation for the S&P 500.

Model	$\alpha = 5\%$					$\alpha = 1\%$					$\alpha = 0.5\%$				
	Loss	**Viol. (%)**	**UC**	**CC**	**IND**	**Loss**	**Viol. (%)**	**UC**	**CC**	**IND**	**Loss**	**Viol. (%)**	**UC**	**CC**	**IND**
HAR-CJ(CPR)	0.133	93 (11.95)	0.000 ***	0.000 ***	0.997	0.041	32 (4.11)	0.000 ***	0.000 ***	0.764	0.025	22 (2.83)	0.000 ***	0.000 ***	0.258
HAR-CJ(ABD)	0.134	91 (11.70)	0.000 ***	0.000 ***	0.873	0.042	30 (3.86)	0.000 ***	0.000 ***	0.876	0.026	21 (2.70)	0.000 ***	0.000 ***	0.280
GJR-GARCH	0.121	58 (7.46)	0.003 ***	0.010 **	0.468	**0.034**	10 (1.29)	0.444	0.654	0.610	**0.019**	**4 (0.51)**	0.956	0.978	0.839
GM-RV	0.119	**52 (6.68)**	0.040 **	0.031 **	0.099 *	0.034	**9 (1.16)**	0.668	0.821	0.646	0.020	5 (0.64)	0.589	0.837	0.799
GM-C(CPR)	0.119	**52 (6.68)**	0.040 **	0.031 **	0.099 *	0.034	**9 (1.16)**	0.668	0.821	0.646	0.019	**4 (0.51)**	0.956	0.978	0.839
GM-CJ(CPR)	0.118	56 (7.20)	0.008 ***	0.025 **	0.563	0.034	**9 (1.16)**	0.668	0.821	0.646	0.019	5 (0.64)	0.589	0.837	0.799
GM-C(ABD)	0.119	**52 (6.68)**	0.040 **	0.031 **	0.099 *	0.034	**9 (1.16)**	0.668	0.821	0.646	0.019	5 (0.64)	0.589	0.837	0.799
GM-CJ(ABD)	**0.118**	55 (7.07)	0.012 **	0.008 ***	0.065 *	0.034	10 (1.29)	0.444	0.654	0.610	0.019	**4 (0.51)**	0.956	0.978	0.839
Realized-GARCH	0.127	80 (10.28)	0.000 ***	0.000 ***	0.958	0.038	25 (3.21)	0.000 ***	0.000 ***	0.197	0.023	14 (1.80)	0.000 ***	0.000 ***	0.473

Notes: Loss is the quantile loss, for which lower values are preferred. Viol. reports the observed number of violations followed by the corresponding violation rate in parentheses. With $T_1 = 778$, the nominal expected counts are 38.90, 7.78, and 3.89 at the 5%, 1%, and 0.5% levels, respectively. Bold identifies the minimum loss or the violation rate closest to the nominal tail probability using unrounded values. UC, CC, and IND are the p-values of the unconditional coverage, conditional coverage, and independence tests, respectively. Shading in loss cells denotes membership in the 90% MCS based on quantile loss. *, **, and *** indicate rejection at the 10%, 5%, and 1% levels, respectively.

At the 1% tail probability, the nominal expected count is 7.78. GM-CJ(CPR) records nine violations, corresponding to a rate of 1.2%, while GM-CJ(ABD) records ten violations, corresponding to 1.3%. Both models remain in the 90% MCS, and the UC, CC, and IND tests do not reject. At the 0.5% tail probability, the expected count is 3.89. GM-CJ(ABD) records four violations and a rate of 0.51%, while GM-CJ(CPR) records five violations and a rate of 0.64%. Both specifications remain in the MCS, and none of the three coverage tests rejects. The deep-tail results therefore combine statistically competitive quantile loss with observed counts close to their nominal expectations under both decompositions.

Across models, HAR-CJ and Realized GARCH produce violation rates well above the nominal probabilities and are repeatedly rejected by the UC and CC tests. At the 1% and 0.5% tail probabilities, the GM-CJ specifications combine MCS membership with non-rejection by the UC, CC, and IND tests.

5.4. Robustness Checks

The robustness analysis uses GM-CJ(CPR) as the primary specification, reflecting its comparatively strong performance in the preceding OOS evaluation, especially for multi-day forecasts and in high-volatility periods. The results for GM-CJ(ABD) are reported alongside the main specification to assess sensitivity to the jump-detection method. Robustness is examined through alternative component allocations, MIDAS block configurations, forecast loss functions, information timing, OOS evaluation windows, and the equity index.

We assess the empirical support for the proposed cross-horizon allocation by comparing GM-CJ with four focused alternatives. GM-RV assigns aggregate realized variance to the long-run component; CJ-L assigns continuous and jump variation separately to the long-run component; GJR-CJ assigns both components to a short-run GJR-GARCH recursion with a constant long-run level; and CJ-LS allows both components to enter both volatility horizons. All specifications use the same data, information convention, rolling window, daily re-estimation, forecast horizons, and realized-variance targets. The comparison uses the 90% MCS with 5000 stationary-bootstrap replications, consistent with the main forecasting evaluation.

Table 10 summarizes the alternative component allocations and their focused MCS results.

Table 10. Alternative component allocations and focused MCS results.

Model	Long-Run Input	Short-Run Input	MCS CPR	MCS ABD	Lowest Loss CPR	Lowest Loss ABD
GM-CJ	C_t	J_t	15/16	16/16	7/16	5/16
GM-RV	Aggregate RV_t	None	10/16	14/16	0/16	7/16
CJ-L	C_t and J_t separately	None	16/16	13/16	9/16	0/16
GJR-CJ	Constant	C_t and J_t	5/16	7/16	0/16	2/16
CJ-LS	C_t and J_t separately	C_t and J_t	9/16	11/16	0/16	2/16

Notes: For each decomposition, the 16 comparisons combine two forecast targets, four horizons, and two loss functions. The MCS columns report the number of comparisons in which each model remains in the 90% model confidence set formed from the five models in the table. The final two columns report the number of comparisons in which each model attains the lowest mean loss.

GM-CJ is retained in 31 of the 32 focused MCSs and records the lowest mean loss in 12 comparisons. CJ-L is also competitive, particularly under CPR, where it attains the lowest loss in nine comparisons. Although jump variation is block-smoothed in CJ-L, its estimated long-run loading is significant under CPR ($\widehat{\lambda}_J = 3.8$, $p = 0.011$) but not under ABD ($\widehat{\lambda}_J = 2.5$, $p = 0.11$). The empirical role of jump variation in the long-run component is therefore sensitive to the jump decomposition. The both-short and both-horizon specifications are retained less often, and their full-sample diagnostics indicate

strong parameter dependence. Appendix B gives the model definitions and the complete cell-level MCS results.

Tables 11–13 organize the robustness results. Table 11 presents the jump-component coefficients and boundary-mixture likelihood-ratio p-values for the CPR and ABD decompositions, together with the conditional-coverage (CC) p-values for GM-CJ(CPR). Because the CC test jointly assesses unconditional coverage and violation independence, the UC and IND results are omitted for concision. Tables 12 and 13 report the percentage loss reductions of GM-CJ(CPR) relative to GM-RV for direct point and average cumulative forecasts, respectively; positive entries favor GM-CJ(CPR). A dagger denotes inclusion of GM-CJ(CPR) in the corresponding 90% MCS formed from all nine models in Table 4. The statistic $\overline{\Delta}_C^Q$ averages the QLIKE improvement over GM-C(CPR) across the four horizons. The final two columns of Tables 12 and 13 count the horizons at which GM-CJ(ABD) is included in the MSE and QLIKE MCSs. Thus, 2/4 denotes inclusion at two of the four horizons.

Replacing the baseline $t-2$ alignment described in Section 3.3 with $t-1$ leaves the coefficient inference unchanged. Under either convention, the CPR and ABD jump-component coefficients are positive, and the boundary-mixture tests reject the null of no jump contribution. For direct point forecasts, the MSE reductions are positive at all four horizons, whereas the QLIKE reductions become positive at $h \geq 5$. The gains are larger for the average cumulative target. At $h = 22$, the MSE and QLIKE reductions are 8.170% and 3.342% under $t-2$ and 7.068% and 3.674% under $t-1$. The mean QLIKE improvement over GM-C(CPR) is positive for both forecast targets under either convention. Across the two timing specifications, GM-CJ(CPR) is included in 15 of the 16 MCSs for direct point forecasts and all 16 MCSs for average cumulative forecasts. For VaR, the CC test under $t-2$ rejects at the 5% tail probability but not at 1% or 0.5%. Under $t-1$, the CC p-values are 0.057, 0.821, and 0.583, implying non-rejection at the 5% significance level for all three tail probabilities.

Forecast gains vary with the OOS window but remain concentrated at multi-day horizons. With an evaluation period beginning in January 2020, GM-CJ(CPR) reduces MSE relative to GM-RV at all four horizons for both forecast targets. The QLIKE reduction for direct point forecasts is -2.568% at $h = 1$ and rises to 2.768%, 6.567%, and 14.963% at $h = 5, 10$, and 22, respectively. For the average cumulative target, the corresponding reductions are -2.568%, 0.812%, 5.603%, and 10.873% at the four levels, respectively. The mean QLIKE improvements over GM-C(CPR) are 4.601% for direct point forecasts and 2.620% for average cumulative forecasts, and GM-CJ(CPR) is included in 14 of the 16 MCSs. The CC test rejects at the 5% tail probability but not at 1% or 0.5%. With an evaluation period beginning in January 2023, MSE reductions for direct point forecasts become negative at $h = 10$ and 22, and the corresponding QLIKE reductions range from -0.250% to 0.103%. The average cumulative QLIKE reductions remain positive for $h \geq 5$. Despite the smaller loss differences, GM-CJ(CPR) is included in 15 of the 16 MCSs, and the CC test does not reject at any of the three tail probabilities. The magnitude of the forecast gain is therefore sample-dependent, but MCS membership and VaR calibration are not tied to the January 2022 starting date.

Table 11. Robustness results for jump-component inference and value-at-risk (VaR).

	CPR		ABD		CC p-Value		
Specification/Sample	$\hat{z}$	**Mixture p**	$\hat{z}$	**Mixture p**	**5%**	**1%**	**0.5%**
S&P 500, baseline ($t-2$)	0.381	0.001	0.325	0.001	0.025	0.821	0.837
S&P 500, timing $t-1$	0.342	0.003	0.274	0.006	0.057	0.821	0.583
S&P 500, OOS start 2020	–	–	–	–	0.004	0.129	0.128
S&P 500, OOS start 2023	–	–	–	–	0.216	0.889	0.959
NASDAQ 100	0.123	0.066	0.160	0.016	0.008	0.052	0.837

Notes: The table reports the jump-component estimate $\hat{z}$ and the boundary-mixture likelihood-ratio p-value under the CPR and ABD decompositions, together with conditional-coverage (CC) p-values for one-day-ahead GM-CJ(CPR) VaR forecasts. Because CC jointly tests unconditional coverage and violation independence, UC and IND results are omitted. The OOS-start rows use the baseline full-sample coefficient estimates; dashes avoid reporting the same estimates repeatedly.

Table 12. Robustness results for direct volatility forecasts.

	MSE				QLIKE				$\overline{\Delta}_C^Q$	ABD MCS	
Specification/Sample	$h=1$	$h=5$	$h=10$	$h=22$	$h=1$	$h=5$	$h=10$	$h=22$	**Mean**	**MSE**	**QLIKE**
S&P 500, baseline ($t-2$)	0.703 †	1.538 †	2.529 †	3.538	−0.719 †	0.295 †	1.289 †	1.395 †	0.788	3/4	2/4
S&P 500, timing $t-1$	0.433 †	1.170 †	2.121 †	3.605 †	−0.243 †	0.361 †	1.449 †	2.300 †	1.191	4/4	4/4
S&P 500, OOS start 2020	6.464 †	4.446 †	1.991 †	2.336 †	−2.568	2.768 †	6.567 †	14.963 †	4.601	4/4	3/4
S&P 500, OOS start 2023	1.057 †	0.979 †	−0.882 †	−1.740	−0.031 †	0.103 †	−0.250 †	0.061 †	−0.063	2/4	4/4
NASDAQ 100	0.437 †	0.400 †	0.713 †	0.932	0.415 †	0.084 †	0.175 †	0.560	0.518	3/4	3/4

Notes: Entries are $\Delta^{s,m}_{RV,h} = 100(L^{s,m}_{GM-RV,h} - L^{s,m}_{GM-CJ(CPR),h})/L^{s,m}_{GM-RV,h}$ for direct point forecasts. Positive values favor GM-CJ(CPR). A superscript † denotes inclusion of GM-CJ(CPR) in the corresponding 90% MCS formed from all nine models. $\overline{\Delta}_C^Q$ averages the four horizon-specific QLIKE improvements over GM-C(CPR). The ABD MCS entries give the number of horizons at which GM-CJ(ABD) belongs to the corresponding MCS.

Table 13. Robustness results for average cumulative volatility forecasts.

	MSE				QLIKE				$\overline{\Delta}_C^Q$	ABD MCS	
Specification/Sample	$h=1$	$h=5$	$h=10$	$h=22$	$h=1$	$h=5$	$h=10$	$h=22$	**Mean**	**MSE**	**QLIKE**
S&P 500, baseline ($t-2$)	0.703 †	2.612 †	4.607 †	8.170 †	−0.719 †	0.312 †	1.476 †	3.342 †	1.671	3/4	3/4
S&P 500, timing $t-1$	0.433 †	1.824 †	3.442 †	7.068 †	−0.243 †	0.501 †	1.435 †	3.674 †	2.051	2/4	4/4
S&P 500, OOS start 2020	6.464 †	6.377 †	5.193 †	4.510 †	−2.568	0.812 †	5.603 †	10.873 †	2.620	4/4	2/4
S&P 500, OOS start 2023	1.057 †	2.432 †	1.215 †	−1.689 †	−0.031 †	0.672 †	0.264 †	1.002 †	0.414	3/4	4/4
NASDAQ 100	0.437 †	1.237 †	1.888 †	2.929 †	0.415 †	0.853 †	1.023 †	0.961 †	1.367	2/4	3/4

Notes: Entries are $\Delta^{s,m}_{RV,h} = 100(L^{s,m}_{GM-RV,h} - L^{s,m}_{GM-CJ(CPR),h})/L^{s,m}_{GM-RV,h}$ for average cumulative forecasts. Positive values favor GM-CJ(CPR). A superscript † denotes inclusion of GM-CJ(CPR) in the corresponding 90% MCS formed from all nine models. $\overline{\Delta}_C^Q$ averages the four horizon-specific QLIKE improvements over GM-C(CPR). The ABD MCS entries give the number of horizons at which GM-CJ(ABD) belongs to the corresponding MCS.

Cross-index robustness is assessed using the NASDAQ 100, a technology-heavy United States equity index. To save space, corresponding robustness results for the Dow Jones Industrial Average (DJI) and Apple Inc. (AAPL) are not tabulated but are available from the authors upon request. The jump-component coefficient is positive under both decompositions. The CPR boundary-mixture test rejects at the 10% level but not at 5% ($p = 0.066$), whereas the ABD test rejects at 5% ($p = 0.016$). All 16 loss reductions relative to GM-RV in Tables 12 and 13 are positive, and GM-CJ(CPR) is included in 14 of the 16 MCSs. For average cumulative forecasts, the MSE reduction increases from 0.437% at $h = 1$ to 2.929% at $h = 22$, and the mean QLIKE improvement over GM-C(CPR) is 1.367%. The CC test rejects at the 5% tail probability but does not reject at 1% ($p = 0.052$) or 0.5%. Hence, the multi-day forecast gains are not confined to the S&P 500, although VaR calibration remains weaker at the 5% tail probability.

Sensitivity to the MIDAS configuration is examined using $(N, K) = (22, 12)$, $(22, 36)$, and $(5, 52)$. Short-run persistence remains between 0.88 and 0.89 under the monthly alternatives and falls to 0.72–0.73 under weekly blocks. The jump coefficient remains positive in every case, although the boundary-mixture test for the weekly ABD specification does not reject at 5% ($p = 0.086$). Under MSE and QLIKE, GM-CJ(CPR) is retained in 15–16 of the 16 full-model and focused MCS comparisons across the four configurations. The corresponding ABD results vary more with the block specification.

Under MAE and RLF, GM-CJ(CPR) records a lower mean loss in 14 of the 16 target–horizon comparisons and is retained in all 16 focused MCSs. The corresponding counts under ABD are six and nine, respectively. The additional criteria therefore preserve the stronger incremental forecasting evidence under CPR and the more configuration-dependent pattern under ABD. Appendix D reports the parameter and forecasting comparisons.

Across the CPR and ABD jump classifications, GM-CJ remains broadly competitive in the focused allocation comparison. Across alternative MIDAS configurations, loss functions, timing conventions, evaluation windows, and the NASDAQ 100 sample, the CPR-based gains remain clearest at $h \geq 5$ and for average cumulative forecasts, with magnitudes that vary across settings. The CC test generally does not reject at the 1% and 0.5% tails.

6. Conclusions

This paper proposes a same-frequency extension of GARCH-MIDAS. Block-smoothed daily continuous variation enters the long-run MIDAS component, whereas demeaned jump variation enters the short-run GJR-GARCH recursion. This allocation is motivated by the different persistence properties of the two realized-volatility components and allows information extracted from high-frequency data to enter conditional variance at the daily frequency. The resulting GARCH-MIDAS-CJ specification retains persistent volatility information through the continuous component while allowing jump variation to affect short-run dynamics.

For the S&P 500, the aligned-sample diagnostics show that continuous variation is substantially more persistent than jump variation under both decompositions, providing the empirical basis for the proposed allocation. Conditional on this allocation, the coefficient and boundary tests support a positive short-run jump effect under CPR and ABD. The focused GM-CJ–GM-C comparison provides more consistent incremental forecasting evidence under CPR and mixed evidence under ABD. In the broader OOS comparison, the clearest gains occur at multi-day horizons and in high-volatility states, whereas the one-day and low-volatility results are more mixed. At the 1% and 0.5% tails, the GM-CJ specifications combine competitive quantile losses with non-rejection of the coverage and independence tests. At the 5% tail, however, violation rates exceed the nominal probability

and exact coverage is rejected, indicating overestimation of the lower-tail return quantile at this level. The CPR-based multi-day pattern remains visible under alternative MIDAS configurations, loss functions, timing conventions, evaluation windows, and the NASDAQ 100 sample, whereas the ABD results remain more configuration-dependent.

These findings have implications for both volatility research and financial risk management. For this application, the results indicate that the predictive content of high-frequency realized measures depends not only on measurement precision but also on how components with different persistence properties are incorporated into conditional variance. Treating continuous and jump variation as distinct inputs provides an interpretable link between realized-volatility measurement and long-run–short-run volatility modeling. For practitioners, the multi-day forecast gains in high-volatility states and the deep-tail VaR calibration results suggest that the proposed allocation may be useful when risk assessments extend beyond one trading day or focus on severe downside outcomes.

Future research may consider signed or state-dependent jump effects, alternative jump-robust estimators, and time-varying jump-component coefficients. Further extensions to multivariate volatility, portfolio risk, and expected shortfall would allow the economic implications of the continuous–jump allocation to be assessed in a broader range of financial applications.

Author Contributions: Conceptualization, M.L. and S.Z.Z.; methodology, M.L.; software, M.L.; validation, M.L. and S.Z.Z.; formal analysis, M.L.; investigation, M.L.; resources, S.Z.Z.; data curation, S.Z.Z.; writing—original draft preparation, M.L.; writing—review and editing, S.Z.Z.; visualization, M.L.; supervision, S.Z.Z.; project administration, S.Z.Z.; funding acquisition, S.Z.Z. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the Guangdong Provincial Key Laboratory of Interdisciplinary Research and Application for Data Science (2022B1212010006).

Data Availability Statement: Restrictions apply to the availability of these data. Data were obtained from the Wind Economic Database and are available at https://www.wind.com.cn/ for users with a valid subscription.

Acknowledgments: During the preparation of this manuscript, the authors used ChatGPT (OpenAI, GPT-5.5) to improve the readability and language quality of the manuscript. The authors have thoroughly reviewed and edited the output and take full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. Well-Definedness of the ABD Variance Allocation

Let

$$c_M = \Phi^{-1}_{1-\alpha_J^*/2} \tag{A1}$$

denote the critical value in the ABD decision rule. We show that $N_t^{\mathrm{ABD}} < M$ under the significance level used in this study.

Suppose, to the contrary, that all M intraday returns are classified as jumps. Equation (14) would then imply

$$|r_{t,i}| > c_M \sqrt{\frac{BV_t}{M}}, \qquad i = 1, \ldots, M. \tag{A2}$$

If $BV_t > 0$, multiplying the inequalities for each pair of adjacent returns and summing over $i = 2, \ldots, M$ gives

$$\sum_{i=2}^{M} |r_{t,i}||r_{t,i-1}| > \frac{M-1}{M} c_M^2 BV_t. \tag{A3}$$

By Equation (5),

$$\sum_{i=2}^{M} |r_{t,i}||r_{t,i-1}| = \mu_1^2 BV_t, \qquad \mu_1^2 = \frac{2}{\pi}. \tag{A4}$$

Consequently, classifying all intervals as jumps would require

$$c_M^2 < \mu_1^2 \frac{M}{M-1}. \tag{A5}$$

For $\alpha_J = 5\%$, the adjusted interval-level significance α_J^* decreases with M, so c_M is nondecreasing in M, whereas $M/(M-1)$ is decreasing. At $M = 2$, c_M^2 is approximately 5.0, while

$$\mu_1^2 \frac{M}{M-1} = \frac{4}{\pi} \approx 1.3. \tag{A6}$$

Thus, Equation (A5) cannot hold for any $M \geq 2$.

If $BV_t = 0$, classifying every interval as a jump would require $|r_{t,i}| > 0$ for every i. All adjacent absolute-return products would then be positive, contradicting Equation (5). Hence, in both cases,

$$N_t^{\mathrm{ABD}} \leq M - 1, \tag{A7}$$

and the denominator in Equation (16) is strictly positive.

Appendix B. Alternative Component Allocations

The allocation exercise supplements the proposed GM-CJ specification and the aggregate-RV benchmark with three variants. For the both-long specification, define the jump block average as

$$\bar{j}_{t-\ell,k}^{(N)} = \frac{1}{N} \sum_{n=1}^{N} \log\left(1 + J_{t-\ell-N(k-1)-n+1}\right), \qquad k = 1, \ldots, K. \tag{A8}$$

Using the same MIDAS weights as the continuous component, the long-run component in CJ-L is

$$\tau_t^L = \exp\left(m + \lambda_C \sum_{k=1}^{K} \phi_k(w_2) \bar{c}_{t-\ell,k}^{(N)} + \lambda_J \sum_{k=1}^{K} \phi_k(w_2) \bar{j}_{t-\ell,k}^{(N)}\right). \tag{A9}$$

CJ-L combines Equation (A9) with the short-run recursion in Equation (25). CJ-LS retains the same long-run component and adds both standardized realized components to the short-run recursion:

$$\begin{aligned} g_t^{LS} = {} & \left(1 - \alpha - \frac{\gamma}{2} - \beta\right) + \left(\alpha + \gamma \mathbf{1}_{\{\varepsilon_{t-1}<0\}}\right) \frac{\varepsilon_{t-1}^2}{\tau_{t-1}^L} + \beta g_{t-1}^{LS} \\ & + \zeta_C \widetilde{C}_{t-1} + \zeta_J \widetilde{J}_{t-1}. \end{aligned} \tag{A10}$$

The both-short GJR-CJ specification uses a constant long-run level and writes conditional variance directly as

$$\sigma_t^2 = \omega + \left(\alpha + \gamma \mathbf{1}_{\{\varepsilon_{t-1}<0\}}\right) \varepsilon_{t-1}^2 + \beta \sigma_{t-1}^2 + \xi_C \widetilde{C}_{t-1} + \xi_J \widetilde{J}_{t-1}. \tag{A11}$$

Here, $\widetilde{C}_t$ and $\widetilde{J}_t$ are standardized using only the observations in the current estimation window. All six decomposition-specific alternatives completed the 778 rolling forecast origins with valid filtered paths and forecasts. The rolling estimates use the reduced $\alpha = 0$ parameterization because the corresponding full-sample estimates place α on its lower boundary.

The full-sample diagnostics distinguish predictive fit from parameter separation. The maximum absolute parameter correlation is approximately 0.99 for both GJR-CJ specifications. For CJ-LS, it is 0.99 under CPR and 0.96 under ABD, with Hessian condition numbers of approximately 1.5×10^{13} and 1.7×10^{10}, respectively. These diagnostics motivate evaluating the alternative placements through their rolling forecast losses.

Table A1. Focused allocation MCS results under the CPR decomposition.

Target	Horizon	Loss	Lowest Loss	Models Retained in the 90% MCS
Point	1	MSE	CJ-L	GM-CJ, CJ-L, GJR-CJ, CJ-LS
Point	1	QLIKE	CJ-L	GM-CJ, GM-RV, CJ-L, GJR-CJ, CJ-LS
Point	5	MSE	CJ-L	GM-CJ, GM-RV, CJ-L, GJR-CJ, CJ-LS
Point	5	QLIKE	CJ-L	GM-CJ, GM-RV, CJ-L
Point	10	MSE	GM-CJ	GM-CJ, CJ-L, CJ-LS
Point	10	QLIKE	GM-CJ	GM-CJ, GM-RV, CJ-L
Point	22	MSE	GM-CJ	GM-CJ, CJ-L, CJ-LS
Point	22	QLIKE	GM-CJ	GM-CJ, GM-RV, CJ-L
Average	1	MSE	CJ-L	GM-CJ, CJ-L, GJR-CJ, CJ-LS
Average	1	QLIKE	CJ-L	GM-CJ, GM-RV, CJ-L, GJR-CJ, CJ-LS
Average	5	MSE	CJ-L	CJ-L
Average	5	QLIKE	CJ-L	GM-CJ, GM-RV, CJ-L
Average	10	MSE	CJ-L	GM-CJ, GM-RV, CJ-L, CJ-LS
Average	10	QLIKE	GM-CJ	GM-CJ, GM-RV, CJ-L
Average	22	MSE	GM-CJ	GM-CJ, CJ-L, CJ-LS
Average	22	QLIKE	GM-CJ	GM-CJ, GM-RV, CJ-L

Notes: Lowest loss identifies the model with the smallest mean loss in each comparison. The MCS is formed from GM-CJ, GM-RV, CJ-L, GJR-CJ, and CJ-LS using 5000 stationary-bootstrap replications.

Table A2. Focused allocation MCS results under the ABD decomposition.

Target	Horizon	Loss	Lowest Loss	Models Retained in the 90% MCS
Point	1	MSE	GJR-CJ	GM-CJ, GM-RV, CJ-L, GJR-CJ, CJ-LS
Point	1	QLIKE	GM-RV	GM-CJ, GM-RV, CJ-L, GJR-CJ, CJ-LS
Point	5	MSE	GM-RV	GM-CJ, GM-RV, CJ-L, GJR-CJ, CJ-LS
Point	5	QLIKE	GM-RV	GM-CJ, GM-RV, CJ-L
Point	10	MSE	GM-CJ	GM-CJ, GM-RV, CJ-LS
Point	10	QLIKE	GM-CJ	GM-CJ, GM-RV, CJ-L
Point	22	MSE	GM-CJ	GM-CJ, CJ-LS
Point	22	QLIKE	GM-CJ	GM-CJ, GM-RV, CJ-L
Average	1	MSE	GJR-CJ	GM-CJ, GM-RV, CJ-L, GJR-CJ, CJ-LS
Average	1	QLIKE	GM-RV	GM-CJ, GM-RV, CJ-L, GJR-CJ, CJ-LS
Average	5	MSE	GM-RV	GM-CJ, GM-RV, CJ-L, GJR-CJ, CJ-LS
Average	5	QLIKE	GM-RV	GM-CJ, GM-RV, CJ-L, CJ-LS
Average	10	MSE	CJ-LS	GM-CJ, GM-RV, CJ-L, GJR-CJ, CJ-LS
Average	10	QLIKE	GM-RV	GM-CJ, GM-RV, CJ-L
Average	22	MSE	CJ-LS	GM-CJ, CJ-LS
Average	22	QLIKE	GM-CJ	GM-CJ, GM-RV, CJ-L

Notes: Lowest loss identifies the model with the smallest mean loss in each comparison. The MCS is formed from GM-CJ, GM-RV, CJ-L, GJR-CJ, and CJ-LS using 5000 stationary-bootstrap replications.

Appendix C. Volatility-State Threshold Sensitivity

Table A3. Sensitivity of high-volatility multi-day loss directions to quantile cutoffs.

Cutoffs	CPR				ABD			
	Avg MSE	Avg QLIKE	Point MSE	Point QLIKE	Avg MSE	Avg QLIKE	Point MSE	Point QLIKE
20/80	3/3	1/3	0/3	2/3	1/3	1/3	0/3	2/3
25/75	3/3	1/3	2/3	2/3	1/3	1/3	1/3	2/3
30/70	3/3	2/3	2/3	2/3	1/3	1/3	2/3	2/3

Notes: Each cell reports the number of $h = 5, 10$, and 22 high-volatility-state comparisons in which GM-CJ has a lower loss than the corresponding GM-C model. The quantile groups are ex post evaluation subsamples based on the corresponding realized out-of-sample target; they are not statistically estimated regimes and do not enter model estimation or forecast formation. No model is re-estimated for this sensitivity analysis.

Appendix D. MIDAS-Configuration and Loss-Function Sensitivity

Table A4 reports full-sample estimates under the submitted configuration and three alternatives. The monthly alternatives vary the maximum lag span while retaining 22 observations per block. The weekly specification uses five observations per block and a maximum span similar to the shorter monthly configuration.

Table A4. Full-sample parameter sensitivity to the MIDAS block configuration.

(N, K)	CPR				ABD			
	Effective Span	Persistence	$\widehat{\zeta}$	p_{LR}	Effective Span	Persistence	$\widehat{\zeta}$	p_{LR}
$(22, 24)$	132	0.89	0.38	0.001	132	0.89	0.33	0.001
$(22, 12)$	110	0.89	0.30	0.005	132	0.89	0.27	0.004
$(22, 36)$	132	0.88	0.43	0.008	132	0.88	0.42	0.002
$(5, 52)$	30	0.73	0.20	0.027	30	0.72	0.13	0.086

Notes: Effective span is the number of trading days covered by the smallest set of most recent blocks containing at least 99.8% of the estimated MIDAS weight. Persistence is $\alpha + \gamma/2 + \beta$. The column p_{LR} reports the boundary-mixture likelihood-ratio p-value for $H_0 : \zeta = 0$. Each configuration uses the full-sample likelihood sample implied by its lag support; likelihood levels are therefore not compared across configurations.

Table A5 summarizes the OOS results under MSE and QLIKE. Each configuration uses the same OOS origins, rolling-window rule, information convention, forecast targets, and MCS procedure.

Table A5. Forecast sensitivity to the MIDAS block configuration.

(N, K)	Full-Model MCS GM-CJ	Focused MCS GM-CJ	Focused MCS GM-C	Lower Mean Loss GM-CJ
Panel A: CPR Decomposition				
$(22, 24)$	15/16	16/16	5/16	14/16
$(22, 12)$	15/16	16/16	8/16	14/16
$(22, 36)$	16/16	15/16	12/16	11/16
$(5, 52)$	15/16	16/16	0/16	16/16
Panel B: ABD Decomposition				
$(22, 24)$	11/16	13/16	13/16	6/16
$(22, 12)$	14/16	12/16	14/16	6/16
$(22, 36)$	15/16	8/16	15/16	6/16
$(5, 52)$	14/16	16/16	4/16	16/16

Notes: The 16 comparisons combine two forecast targets, four horizons, and two loss functions. The full-model MCS contains the nine models in Table 4. The focused MCS contains GM-CJ and its nested GM-C specification under the same decomposition. The final column is descriptive and counts comparisons in which GM-CJ has the lower mean loss. All MCSs use a 90% confidence level and 5000 stationary-bootstrap replications.

Table A6 reports the percentage loss reduction of GM-CJ relative to GM-C under MAE and RLF.

Table A6. Additional-loss comparison of GM-CJ and GM-C.

Decomposition	Target	Loss	$h = 1$	$h = 5$	$h = 10$	$h = 22$
CPR	Average	MAE	1.12	1.89	2.38	3.74
CPR	Point	MAE	1.12	1.48	2.28	2.21
CPR	Average	RLF	−0.28	1.41	3.69	6.19
CPR	Point	RLF	−0.28	1.28	2.45	2.57
ABD	Average	MAE	−2.07	−2.22	−1.72	0.90
ABD	Point	MAE	−2.07	−0.56	0.63	1.79
ABD	Average	RLF	−3.41	−3.44	−1.28	2.88
ABD	Point	RLF	−3.41	−0.93	1.08	2.63

Notes: Entries are percentage loss reductions, $100(L_{\mathrm{GM-C}} - L_{\mathrm{GM-CJ}})/L_{\mathrm{GM-C}}$; positive values favor GM-CJ. Shaded cells indicate that GM-CJ remains in the corresponding 90% focused MCS. Across MAE and RLF, GM-CJ is retained in 16/16 focused MCSs under CPR and 9/16 under ABD. In the corresponding nine-model MCSs, it is retained in 12/16 and 6/16 comparisons, respectively. All MCSs use 5000 stationary-bootstrap replications.

References

1. Engle, R.F.; Patton, A.J. What good is a volatility model? *Quant. Financ.* **2001**, *1*, 237–245. [CrossRef]
2. Engle, R.F. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica* **1982**, *50*, 987–1007. [CrossRef]
3. Bollerslev, T. Generalized autoregressive conditional heteroskedasticity. *J. Econom.* **1986**, *31*, 307–327. [CrossRef]
4. Nelson, D.B. Conditional heteroskedasticity in asset returns: A new approach. *Econometrica* **1991**, *59*, 347–370. [CrossRef]
5. Glosten, L.R.; Jagannathan, R.; Runkle, D.E. On the relation between the expected value and the volatility of the nominal excess return on stocks. *J. Financ.* **1993**, *48*, 1779–1801. [CrossRef]
6. Engle, R.F.; Lee, G.G.J. A long-run and short-run component model of stock return volatility. In *Cointegration, Causality, and Forecasting: A Festschrift in Honour of Clive W. J. Granger*; Engle, R.F., White, H., Eds.; Oxford University Press: Oxford, UK, 1999; pp. 475–497.
7. Baillie, R.T.; Bollerslev, T.; Mikkelsen, H.O. Fractionally integrated generalized autoregressive conditional heteroskedasticity. *J. Econom.* **1996**, *74*, 3–30. [CrossRef]
8. Andersen, T.G.; Bollerslev, T. Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *Int. Econ. Rev.* **1998**, *39*, 885–905. [CrossRef]
9. Andersen, T.G.; Bollerslev, T.; Diebold, F.X.; Ebens, H. The distribution of realized stock return volatility. *J. Financ. Econ.* **2001**, *61*, 43–76. [CrossRef]
10. Barndorff-Nielsen, O.E.; Shephard, N. Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *J. R. Stat. Soc. Ser. B (Stat. Methodol.)* **2002**, *64*, 253–280. [CrossRef]
11. Andersen, T.G.; Bollerslev, T.; Diebold, F.X.; Labys, P. Modeling and forecasting realized volatility. *Econometrica* **2003**, *71*, 579–625. [CrossRef]
12. Corsi, F. A simple approximate long-memory model of realized volatility. *J. Financ. Econom.* **2009**, *7*, 174–196. [CrossRef]
13. Hansen, P.R.; Huang, Z.; Shek, H.H. Realized GARCH: A joint model for returns and realized measures of volatility. *J. Appl. Econom.* **2012**, *27*, 877–906. [CrossRef]
14. Shephard, N.; Sheppard, K. Realising the future: Forecasting with high-frequency-based volatility (HEAVY) models. *J. Appl. Econom.* **2010**, *25*, 197–231. [CrossRef]
15. Hansen, P.R.; Lunde, A. Realized variance and market microstructure noise. *J. Bus. Econ. Stat.* **2006**, *24*, 127–161. [CrossRef]
16. McAleer, M.; Medeiros, M.C. Realized volatility: A review. *Econom. Rev.* **2008**, *27*, 10–45. [CrossRef]
17. Engle, R.F.; Ghysels, E.; Sohn, B. Stock market volatility and macroeconomic fundamentals. *Rev. Econ. Stat.* **2013**, *95*, 776–797. [CrossRef]
18. Asgharian, H.; Hou, A.J.; Javed, F. The importance of the macroeconomic variables in forecasting stock return variance: A GARCH-MIDAS approach. *J. Forecast.* **2013**, *32*, 600–612. [CrossRef]
19. Conrad, C.; Loch, K. Anticipating long-term stock market volatility. *J. Appl. Econom.* **2015**, *30*, 1090–1114. [CrossRef]
20. Conrad, C.; Kleen, O. Two are better than one: Volatility forecasting using multiplicative component GARCH-MIDAS models. *J. Appl. Econom.* **2020**, *35*, 19–45. [CrossRef]

21. Engle, R.F.; Rangel, J.G. The spline-GARCH model for low-frequency volatility and its global macroeconomic causes. *Rev. Financ. Stud.* **2008**, *21*, 1187–1222. [CrossRef]
22. Borup, D.; Jakobsen, J.S. Capturing volatility persistence: A dynamically complete realized EGARCH-MIDAS model. *Quant. Financ.* **2019**, *19*, 1839–1855. [CrossRef]
23. Huang, X.; Tauchen, G. The relative contribution of jumps to total price variance. *J. Financ. Econom.* **2005**, *3*, 456–499. [CrossRef]
24. Andersen, T.G.; Bollerslev, T.; Diebold, F.X. Roughing it up: Including jump components in the measurement, modeling, and forecasting of return volatility. *Rev. Econ. Stat.* **2007**, *89*, 701–720. [CrossRef]
25. Patton, A.J.; Sheppard, K. Good volatility, bad volatility: Signed jumps and the persistence of volatility. *Rev. Econ. Stat.* **2015**, *97*, 683–697. [CrossRef]
26. Corsi, F.; Pirino, D.; Renò, R. Threshold bipower variation and the impact of jumps on volatility forecasting. *J. Econom.* **2010**, *159*, 276–288. [CrossRef]
27. Andersen, T.G.; Bollerslev, T.; Dobrev, D. No-arbitrage semi-martingale restrictions for continuous-time volatility models subject to leverage effects, jumps and i.i.d. noise: Theory and testable distributional implications. *J. Econom.* **2007**, *138*, 125–180. [CrossRef]
28. Papantonis, I.; Rompolis, L.; Tzavalis, E. Improving variance forecasts: The role of Realized Variance features. *Int. J. Forecast.* **2023**, *39*, 1221–1237. [CrossRef]
29. Ghysels, E.; Sinko, A.; Valkanov, R. MIDAS regressions: Further results and new directions. *Econom. Rev.* **2007**, *26*, 53–90. [CrossRef]
30. Amendola, A.; Candila, V.; Gallo, G.M. Choosing the frequency of volatility components within the Double Asymmetric GARCH–MIDAS–X model. *Econom. Stat.* **2021**, *20*, 12–28. [CrossRef]
31. Wang, L.; Zhao, C.; Liang, C.; Jiu, S. Predicting the volatility of China's new energy stock market: Deep insight from the realized EGARCH-MIDAS model. *Financ. Res. Lett.* **2022**, *48*, 102981. [CrossRef]
32. Wu, X.; Zhao, A.; Cheng, T. A Real-Time GARCH-MIDAS model. *Financ. Res. Lett.* **2023**, *56*, 104103. [CrossRef]
33. Segnon, M.; Gupta, R.; Wilfling, B. Forecasting stock market volatility with regime-switching GARCH-MIDAS: The role of geopolitical risks. *Int. J. Forecast.* **2024**, *40*, 29–43. [CrossRef]
34. Conrad, C.; Engle, R.F. Modelling Volatility Cycles: The MF2-GARCH Model. *J. Appl. Econom.* **2025**, *40*, 438–454. [CrossRef]
35. Barndorff-Nielsen, O.E.; Shephard, N. Power and bipower variation with stochastic volatility and jumps. *J. Financ. Econom.* **2004**, *2*, 1–37. [CrossRef]
36. Barndorff-Nielsen, O.E.; Shephard, N. Econometrics of testing for jumps in financial economics using bipower variation. *J. Financ. Econom.* **2006**, *4*, 1–30. [CrossRef]
37. Lee, S.S.; Mykland, P.A. Jumps in financial markets: A new nonparametric test and jump dynamics. *Rev. Financ. Stud.* **2008**, *21*, 2535–2563. [CrossRef]
38. Aït-Sahalia, Y.; Jacod, J. Testing for jumps in a discretely observed process. *Ann. Stat.* **2009**, *37*, 184–222. [CrossRef]
39. Andersen, T.G.; Dobrev, D.; Schaumburg, E. Jump-robust volatility estimation using nearest neighbor truncation. *J. Econom.* **2012**, *169*, 75–93. [CrossRef]
40. Caporin, M. The Role of Jumps in Realized Volatility Modeling and Forecasting. *J. Financ. Econom.* **2023**, *21*, 1143–1168. [CrossRef]
41. Bu, R.; Hizmeri, R.; Izzeldin, M.; Murphy, A.; Tsionas, M. The contribution of jump signs and activity to forecasting stock price volatility. *J. Empir. Financ.* **2023**, *70*, 144–164. [CrossRef]
42. Caporin, M.; Di Fonzo, T.; Girolimetto, D. Exploiting Intraday Decompositions in Realized Volatility Forecasting: A Forecast Reconciliation Approach. *J. Financ. Econom.* **2024**, *22*, 1759–1784. [CrossRef]
43. Al Rababaa, A.R.; Mensi, W.; McMillan, D.; Kang, S.H. Forecasting the Realized Volatility of Stock Markets: The Roles of Jumps and Asymmetric Spillovers. *J. Forecast.* **2025**, *44*, 1294–1325. [CrossRef]
44. Audrino, F.; Hu, Y. Volatility Forecasting: Downside Risk, Jumps and Leverage Effect. *Econometrics* **2016**, *4*, 8. [CrossRef]
45. Conrad, C.; Loch, K.; Rittler, D. On the macroeconomic determinants of long-term volatilities and correlations in U.S. stock and crude oil markets. *J. Empir. Financ.* **2014**, *29*, 26–40. [CrossRef]
46. Bougerol, P.; Picard, N. Strict Stationarity of Generalized Autoregressive Processes. *Ann. Probab.* **1992**, *20*, 1714–1730. [CrossRef]
47. Straumann, D.; Mikosch, T. Quasi-Maximum-Likelihood Estimation in Conditionally Heteroscedastic Time Series: A Stochastic Recurrence Equations Approach. *Ann. Stat.* **2006**, *34*, 2449–2495. [CrossRef]
48. Bollerslev, T.; Wooldridge, J.M. Quasi-Maximum Likelihood Estimation and Inference in Dynamic Models with Time-Varying Covariances. *Econom. Rev.* **1992**, *11*, 143–172. [CrossRef]
49. Self, S.G.; Liang, K.Y. Asymptotic Properties of Maximum Likelihood Estimators and Likelihood Ratio Tests under Nonstandard Conditions. *J. Am. Stat. Assoc.* **1987**, *82*, 605–610. [CrossRef]
50. Andrews, D.W.K. Testing When a Parameter Is on the Boundary of the Maintained Hypothesis. *Econometrica* **2001**, *69*, 683–734. [CrossRef]

51. Liu, L.Y.; Patton, A.J.; Sheppard, K. Does anything beat 5-minute RV? A comparison of realized measures across multiple asset classes. *J. Econom.* **2015**, *187*, 293–311. [CrossRef]
52. Francq, C.; Zakoïan, J.M. Quasi-Maximum Likelihood Estimation in GARCH Processes When Some Coefficients Are Equal to Zero. *Stoch. Process. Their Appl.* **2007**, *117*, 1265–1284. [CrossRef]
53. Hansen, P.R.; Lunde, A. A forecast comparison of volatility models: Does anything beat a GARCH(1,1)? *J. Appl. Econom.* **2005**, *20*, 873–889. [CrossRef]
54. Patton, A.J. Volatility forecast comparison using imperfect volatility proxies. *J. Econom.* **2011**, *160*, 246–256. [CrossRef]
55. Moroke, N.; Makatjane, K.; Shoko, C. Evaluating Forecast Accuracy: A Comparative Error-Based Analysis of ANN and Markov-Switching Models in Inflation Prediction. *J. Stat. Appl. Probab.* **2026**, *15*, 579–597. [CrossRef]
56. Hansen, P.R.; Lunde, A.; Nason, J.M. The model confidence set. *Econometrica* **2011**, *79*, 453–497. [CrossRef]
57. González-Rivera, G.; Lee, T.H.; Mishra, S. Forecasting volatility: A reality check based on option pricing, utility function, value-at-risk, and predictive likelihood. *Int. J. Forecast.* **2004**, *20*, 629–645. [CrossRef]
58. Bernardi, M.; Catania, L. Comparison of Value-at-Risk models using the MCS approach. *Comput. Stat.* **2016**, *31*, 579–608. [CrossRef]
59. Bernardi, M.; Catania, L.; Petrella, L. Are news important to predict the Value-at-Risk? *Eur. J. Financ.* **2017**, *23*, 535–572. [CrossRef]
60. Kupiec, P.H. Techniques for verifying the accuracy of risk measurement models. *J. Deriv.* **1995**, *3*, 73–84. [CrossRef]
61. Christoffersen, P.F. Evaluating interval forecasts. *Int. Econ. Rev.* **1998**, *39*, 841–862. [CrossRef]
62. Christoffersen, P.; Hahn, J.; Inoue, A. Testing and comparing value-at-risk measures. *J. Empir. Financ.* **2001**, *8*, 325–342. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.