Volatility Forecasting Factors

Marco Cinquetti
Lancaster University Management School, Lancaster, UK
and
Seok Young Hong
Nanyang Technological University, Singapore
and
Ingmar Nolte
Lancaster University Management School, Lancaster, UK
and
Sandra Nolte
Lancaster University Management School, Lancaster, UK

Abstract

This paper examines the predictive role of high-frequency factor volatilities in modeling the volatility of individual stocks. We develop a dynamic forecasting framework that selects the most informative factor-specific realized volatility from a large cross-section of asset pricing anomalies. Embedded in a log-linear specification, the model integrates both market-wide and idiosyncratic components, allowing for a flexible representation of volatility dynamics. Empirical results based on a broad universe of U.S. equities demonstrate that the proposed method significantly outperforms standard benchmarks, both statistically and economically. The findings underscore the importance of incorporating high-frequency cross-sectional information in volatility modeling, offering a scalable and interpretable approach to understanding time-varying risks in equity markets.

Keywords: Factor models; High-frequency data; Realized volatility

1 Introduction

Volatility plays a fundamental role in finance as it underlies the risk-reward dynamics defining modern financial theories, from investment decision-making to monetary policies. Despite its centrality, the modeling of volatility remains largely uninformed about the multivariate frameworks that dominate empirical asset pricing. From the first factor model of Sharpe (1964) to the "zoo" of factors in Harvey, Liu & Zhu (2016), the literature on financial returns has long embraced systematic, cross-sectional structures; in contrast, volatility has mainly been modeled as an idiosyncratic, asset-specific process. As Bollerslev (2022) notes, volatility modeling has been slow to internalize the inherently multivariate nature of market dynamics, often ignoring the factor structures that drive both return co-movements and correlated risk exposures. The result is a methodological gap: whereas expected returns are modeled via linear exposures to observable or latent factors, volatility is typically forecasted without acknowledging systematic drivers.

Any volatility modeling framework ultimately hinges on how precisely variance can be measured. Early contributions such as the ARCH model of Engle (1982) and its generalized version by Bollerslev (1986) introduced conditional heteroskedasticity as a time-varying property of return series. Parallel to these, the stochastic volatility models by Taylor (1982) offered a latent process formulation. The introduction of high-frequency financial data marked a fundamental shift, enabling, under ideal sampling conditions, nonparametric estimation of risk with realized volatility (RV) measures. However, the discrete and discontinuous nature of financial markets poses challenges for the unbiased and consistent estimation of realized volatility: multi-scale estimators (Zhang et al. 2005), least-squares noise-robust inference methods (Nolte & Voev 2012), and pre-averaging methods (Christensen et al. 2014) address market microstructure noise; bipower and multipower estimators (Barndorff-Nielsen & Shephard 2004), truncated estimators (Mancini 2009) and some combinations thereof (Corsi et al. 2010) focus on jumps; recent contributions

(Andersen et al. 2023, Li et al. 2025) provide robust estimation in presence of short-lived extreme price movements. These advancements have made it possible to recover the latent volatility process efficiently, providing a reliable instrument for volatility forecasting models.

One of the most influential frameworks to leverage realized volatility is the Heterogeneous Autoregressive model of Corsi (2009). The HAR model parsimoniously captures the long-memory behavior of volatility by including lagged daily, weekly, and monthly realized volatilities as regressors, approximating the effect of a broader lag distribution in a simple linear specification. Volatility forecasting has further improved with extensions such as the semivariance HAR (SHAR) model of Patton & Sheppard (2015), allowing for conditional asymmetries through signed returns, or the HARQ of Bollerslev, Patton & Quaedvlieg (2016), introducing realized quarticity to account for temporal variations in RV measurement errors. Despite these refinements, a foundational limitation remains: HAR models are inherently univariate, thus agnostic to any systematic risk factors.

The limitations of univariate volatility models have motivated a shift towards specifications incorporating cross-sectional information. A prominent example is the market-HAR model introduced by Hizmeri, Izzeldin, Nolte & Pappas (2022), which augments the standard framework with market-level realized (co)variances and semi(co)variances. Building on similar reasoning, the multiplicative volatility factor (MVF) model proposed by Ding, Engle, Li & Zheng (2025) presents a parsimonious structure in which the realized variance of each stock is expressed as the product of a latent common volatility factor and an idiosyncratic residual. Empirically, both models demonstrate significant gains in forecast accuracy, confirming a stylized fact: stock volatilities co-move over time, often driven by aggregated shocks rather than isolated firm-level events. This insight is well-grounded in the empirical literature. Early studies by Engle, Ito & Lin (1988) and Calvet, Fisher & Thompson (2006) document volatility transmission and comovement across markets and horizons, from intraday spillovers across FX trading

centers to multifrequency comovement across exchange rates. Complementing this, Herskovic, Kelly, Lustig & Van Nieuwerburgh (2016) identify a latent common idiosyncratic volatility factor that explains a significant fraction of cross-sectional volatility dispersion. More recent works emphasize the role of firm-level linkages: Herskovic, Kelly, Lustig & Van Nieuwerburgh (2020) show that firms embedded in central positions within economic networks exhibit stronger volatility co-movement. At the macro level, Bollerslev, Hood, Huss & Pedersen (2018) and Engle & Campos-Martins (2023) provide evidence of global volatility factors driving fluctuations in equity markets, reinforcing the importance of modeling volatility beyond the firm-specific scale.

The observed heterogeneity in volatility patterns extends beyond a single driver. Figure 1 displays the annualized realized volatility of the high-frequency market (MKT) factor against the other four Fama & French (2015) factors, namely size (SMB), value (HML), profitability (RMW), and investment (CMA). Every time series exhibits distinct and persistent dynamics over time, suggesting that volatility may originate from multiple, structurally distinct sources. Recent works by Barigozzi & Hallin (2017) and Kapadia, Linn & Paye (2024) show that volatility in large panels of factor returns admits a low-dimensional structure, with a pervasive common component alongside factor-specific dynamics. In this context, the identification of the most informative factor variance must be flexible, thereby motivating adaptive multi-factor frameworks that update their forecasting structure in response to shifting volatility regimes.

The present paper introduces a novel volatility model that extends the HAR structure by dynamically selecting the most informative components among a large cross-section of high-frequency asset pricing factors. Specifically, we construct daily realized variances based on second-level intraday returns for 287 observable anomalies, covering the full set proposed in Fama & French (2018), Chen & Zimmermann (2022) and Jensen, Kelly & Pedersen (2023). Methodologically, this approach contributes to high-dimensional forecasting by imposing

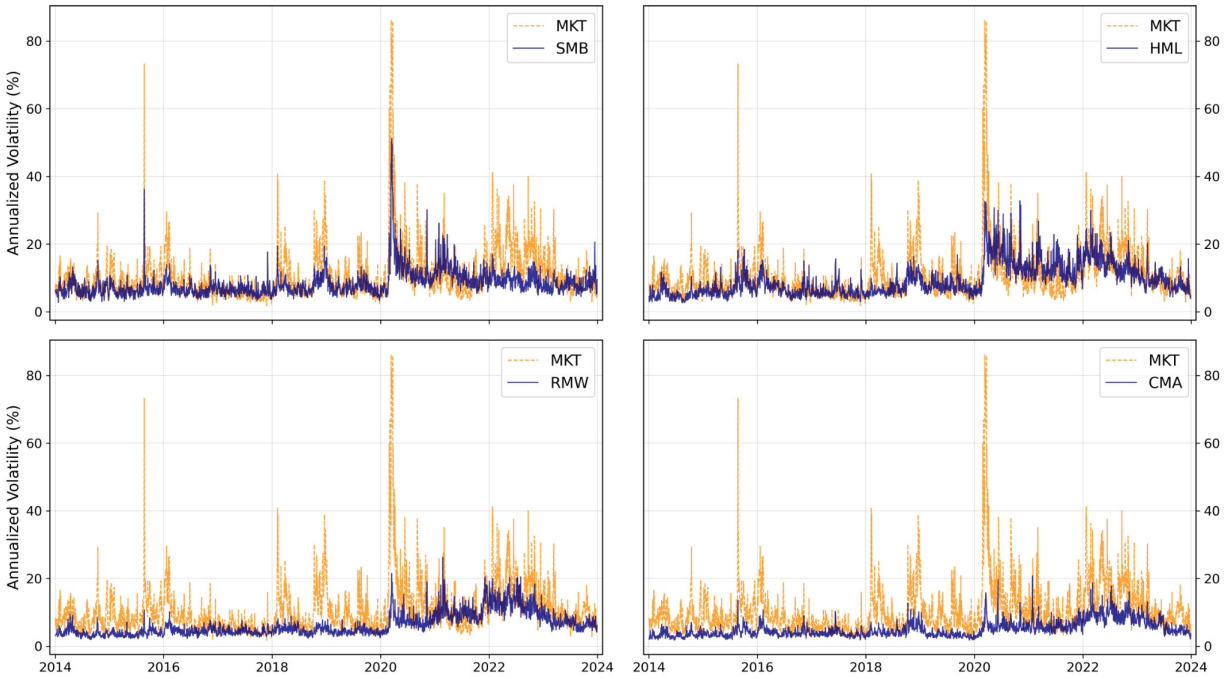

Figure 1: Annualized 5-minute realized volatility of Fama-French five factors (MKT, SMB, HML, RMW, CMA) from 2014 to 2023. The figure illustrates the heterogeneous dynamics of factor volatilities.

structural sparsity through economically motivated selection. It differs from penalized-regression return-forecasting approaches that use LASSO-style selection in high-frequency prediction settings (Chinco et al. 2019) and from broader machine-learning forecasting toolkits that trade off flexibility and interpretability across nonlinear learners (Gu et al. 2020). It is also distinct from nonparametric characteristic-selection frameworks that combine adaptive group-LASSO selection with flexible estimation of predictor effects (Freyberger et al. 2020). Moreover, it complements dynamic-factor approaches that decompose large panels of volatilities and deliver predictive objects such as conditional prediction intervals (Barigozzi & Hallin 2020, Luciani & Veredas 2015), high-dimensional integrated covariance estimation (Zheng & Li 2011), and large-system stochastic volatility methods for tail-risk measures (Asai et al. 2015) as well as factor-based realized volatility forecasting that accounts for jumps (Atak & Kapetanios 2013).

We evaluate forecasting performance across more than one thousand U.S. equities, markedly exceeding the cross-sectional coverage of typical studies in the volatility

forecasting literature. The proposed model is benchmarked against several prominent volatility forecasting frameworks, delivering statistically significant improvements in forecast accuracy. Assets inherit volatility from common drivers: shocks to factor-level uncertainty propagate to individual variance through time-varying exposures, so realized factor volatility at multiple horizons add information beyond an asset's own history. The robustness of these gains is confirmed across multiple realized volatility estimators, forecasting horizons, and loss functions. These results establish the necessity of incorporating factor-specific volatility components in forecasting models, and show that informational efficiency depends not only on the precision of realized variance measurement but also on the structural selection of its predictors.

The remainder of the paper is structured as follows. Section 2 outlines the econometric framework, model specification, and theoretical foundations of the log-HAR framework with adaptive factor volatilities. We prove the strong consistency of our selection procedure, where we show the quasi-likelihood (QLIKE)-based selection rule asymptotically identifies the factor that truly drives volatility. We also characterize how the mis-selection probability responds to measurement error. Section 3 describes the dataset construction and implementation details. Section 4 presents the forecasting results and factor analysis. Section 5 concludes.

2 Theoretical framework

2.1 Setup and background theory

Throughout the paper, we work on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$ with a complete and right-continuous filtration. For each asset $i = 1, \ldots, N$, the log-price process $(P_{i,t})_{t\geq 0}$ is defined on this space, adapted to $(\mathcal{F}_t)$, and is an Itô semimartingale of the form

$$dP_{i,t} = \mu_{i,t}\, dt + \sigma_{i,t}\, dW_{i,t}; \qquad t \geq 0, \tag{1}$$

where the drift $\mu_{i,t}$ and instantaneous volatility $\sigma_{i,t}$ are $(\mathcal{F}_t)$-progressively measurable processes such that, for every $T < \infty$, $\int_0^T |\mu_{i,t}|dt < \infty$ a.s. and $\int_0^T \sigma_{i,t}^2 dt < \infty$ a.s., and $W_{i,t}$ is a standard $(\mathcal{F}_t)$-Brownian motion.

Writing the trading day $t \geq 1$ as the unit-length interval $(t-1, t]$, we define the latent integrated variance for asset i on day t as

$$IV_{i,t} := \int_{t-1}^{t} \sigma_{i,s}^2 \, ds. \tag{2}$$

The integrated variance aggregates the entire intraday volatility path, and is the quantity we ultimately seek to forecast at daily and longer horizons. To maintain conciseness in the exposition, we abstract from discontinuous price moves here; jumps and other extreme price movements are incorporated through robust estimators in the empirical analysis presented in Section 4.

Although the integrated variance $IV_{i,t}$ is not observable, it is well known that high-frequency returns yield a consistent nonparametric estimator. Specifically, let $\Delta = 1/M$ be the intraday sampling interval and index observations by $j = 1, \ldots, M$. On the equally spaced intraday grid, the j-th intraday return is $r_{i,t,j} = P_{i,t-1+j\Delta} - P_{i,t-1+(j-1)\Delta}$. Under the continuous semimartingale assumption and in the absence of market microstructure noise, the realized variance

$$RV_{i,t} := \sum_{j=1}^{M} r_{i,t,j}^2 = \sum_{j=1}^{M} \left(P_{i,t-1+j\Delta} - P_{i,t-1+(j-1)\Delta} \right)^2 \tag{3}$$

converges in probability to $IV_{i,t}$ as the mesh $\Delta \to 0$ (i.e. $M \to \infty$); the analogous joint convergence holds in multivariate sense for the realized covariance matrix, and hence for any fixed linear portfolio. If there was a jump component in (1), then (3) would converge in probability to the quadratic variation, which incorporates the jump variations in addition to the integrated variance. See Theorems 4.47 and 4.52 of Jacod & Shiryaev (2003) for details. The limiting distribution of the realized variance has been widely investigated in the literature under diverse conditions, see Aït-Sahalia & Jacod (2014) for a comprehensive

exposition.

We construct asset pricing factors from the same universe of stocks as the individual assets. Let $r_{k,t}$ denote the daily log return on an asset pricing factor $k \in \mathcal{K}$, with $\mathcal{K}$ denoting the finite index set corresponding to the set of candidate factors (e.g., MKT, SMB, HML).[1] Writing $r_{k,t,j}$ for the j-th intraday log return of factor k computed on the equally spaced grid $\{t-1+\Delta, t-1+2\Delta, \ldots, t\}$ with mesh $\Delta = 1/M$, we estimate the factor volatility nonparametrically by the realized variance

$$FRV_{k,t} := \sum_{j=1}^{M} r_{k,t,j}^2, \tag{4}$$

the *realized factor variance*. By continuous mapping, the consistency and limiting distribution results extend to the realized variance of the factor return, provided the sampling is free of microstructure noise and the factor is implemented as a self-financing, tradable linear portfolio with weights that are $\mathcal{F}_{t-1}$-measurable and of bounded variation on $(t-1, t]$.

The asset- and factor-level realized variances defined above constitute the observable building blocks for the theoretical framework developed in the next subsections.

2.2 Adaptive Factor-Driven Volatility Models

We propose a volatility forecasting framework that augments the heterogeneous autoregressive (HAR) structure with multiplicative factor components. To ensure the non-negativity of the variance, to alleviate right-skewness and heavy tails, and to render the multiplicative decomposition additive, we adopt a log-linear specification and work with $\log RV_{i,t}$, i.e., the log realized variance of stock i on day t.

We model $\log RV_{i,t}$ as the linear combination of three log predictor blocks. First, the market proxy, which is defined as the cross-sectional average of realized variances across all

[1] All factors examined are diversified in the sense of Ross (2013), so the associated portfolios bear negligible firm-specific risk. This feature distinguishes our analysis from Herskovic, Kelly, Lustig & Van Nieuwerburgh (2016) and related studies that focus on commonality in the volatility of firm-specific returns.

stocks

$$CRV_t := \frac{1}{N}\sum_{i=1}^{N} RV_{i,t}, \tag{5}$$

hereafter referred to as *common realized variance* (CRV). Second, the stock-specific component

$$\xi_{i,t} := \frac{RV_{i,t}}{CRV_t}, \tag{6}$$

defined as the multiplicative residual of stock i's variance relative to the cross-section. Third, the factor specific component $FRV_{k^\star_{i,t}}$, which is the *realized factor variance* (4) corresponding to the high-frequency factor $k^\star_{i,t} \in \mathcal{K}$, where $k^\star_{i,t}$ is chosen adaptively as the factor whose realized variance lags yield the best in-sample explanatory power for stock i.[2] See (8) below and the discussion that follows.

Each component adds to the model daily, weekly, and monthly lags. Specifically, for each stock i, day t and forecasting horizon h, the log-realized variance $\log RV_{i,t+h}$ follows the linear model

$$\begin{aligned}
\log RV_{i,t+h} &= \beta_0 + \beta^{(d)}_{CRV}\log CRV^{(d)}_t + \beta^{(w)}_{CRV}\log CRV^{(w)}_t + \beta^{(m)}_{CRV}\log CRV^{(m)}_t \\
&\quad + \beta^{(d)}_{\xi}\log \xi^{(d)}_{i,t} + \beta^{(w)}_{\xi}\log \xi^{(w)}_{i,t} + \beta^{(m)}_{\xi}\log \xi^{(m)}_{i,t} \\
&\quad + \beta^{(d)}_{k^\star}\log FRV^{(d)}_{k^\star_{i,t}} + \beta^{(w)}_{k^\star}\log FRV^{(w)}_{k^\star_{i,t}} + \beta^{(m)}_{k^\star}\log FRV^{(m)}_{k^\star_{i,t}} + \varepsilon_{i,t+h} \\
&= \beta_0 + \beta^{\top}_{\text{CRV}} \underbrace{\begin{bmatrix} \log CRV^{(d)}_t \\ \log CRV^{(w)}_t \\ \log CRV^{(m)}_t \end{bmatrix}}_{\text{Market/Common block}} + \beta^{\top}_{\xi} \underbrace{\begin{bmatrix} \log \xi^{(d)}_{i,t} \\ \log \xi^{(w)}_{i,t} \\ \log \xi^{(m)}_{i,t} \end{bmatrix}}_{\text{Stock block}} + \beta^{\top}_{k^\star} \underbrace{\begin{bmatrix} \log FRV^{(d)}_{k^\star_{i,t}} \\ \log FRV^{(w)}_{k^\star_{i,t}} \\ \log FRV^{(m)}_{k^\star_{i,t}} \end{bmatrix}}_{\text{Factor block}} + \varepsilon_{i,t+h},
\end{aligned} \tag{7}$$

where the coefficients $\beta_{CRV} = (\beta^{(d)}_{CRV}, \beta^{(w)}_{CRV}, \beta^{(m)}_{CRV})$, $\beta_\xi = (\beta^{(d)}_{\xi}, \beta^{(w)}_{\xi}, \beta^{(m)}_{\xi})$ and $\beta_{k^\star} = (\beta^{(d)}_{k^\star}, \beta^{(w)}_{k^\star}, \beta^{(m)}_{k^\star})$ are estimated by the OLS over a rolling estimation window of fixed length L. The superscripts (d), (w), (m) refer to the daily ($x^{(d)}_t = x_t$), weekly ($x^{(w)}_t = \frac{1}{5}\sum_{\tau=0}^{4} x_{t-\tau}$),

[2] Equation (4) defines $FRV_{k,t}$ as the realized variance of factor k at date t. When the selected factor index is $k^\star_{i,t}$ (chosen for stock i at date t), the corresponding series is $FRV_{k^\star_{i,t},t}$. To avoid reporting the same time index, we henceforth write $FRV_{k^\star_{i,t}}$

and monthly ($x_t^{(m)} = \frac{1}{22}\sum_{\tau=0}^{21} x_{t-\tau}$) averages, respectively, for $x_t = CRV_t$, $\xi_{i,t}$, or $FRV_{k_{i,t}^\star}$.

For a selection window of length S, the high-frequency factor to be chosen by $k_{i,t}^\star$ is the one delivering the *best in-sample quasi-likelihood* (QLIKE) performance. That is, we choose the factor that minimizes the QLIKE loss

$$k_{i,t}^\star := \arg\min_{k\in\mathcal{K}} \left\{ \frac{1}{S} \sum_{s=t-S+1}^{t} \left[\log\left(\frac{\widehat{RV}_{i,s}^{\,k}}{RV_{i,s}} \right) + \frac{RV_{i,s}}{\widehat{RV}_{i,s}^{\,k}} - 1 \right] \right\}, \tag{8}$$

where $RV_{i,s}$ is the measured realized variance and $\widehat{RV}_{i,s}^{\,k}$ is the forecasted (with candidate factor k) realized variance of stock i on day s of the selection period.[3] The quantity $\widehat{RV}_{i,s}^{\,k}$ is simply the forecast for $RV_{i,s}$ produced by the proposed model when the factor block is built from candidate factor k (and the remaining CRV and ξ blocks are unchanged), with parameters estimated on the rolling estimation window.

Remark. The estimation window L and the selection window S serve different purposes. The former is used to estimate the coefficients in (7) while the latter is used to score candidate factors in (8). In our specification, we impose $S \leq L$ so that the S evaluation observations lie within the L-day estimation sample, ensuring that the QLIKE criterion is computed in-sample at t under common (fixed) parameter estimates.

The final specification (7) preserves the interpretability of the standard HAR structure while expanding its information set to include an extensive set of high-frequency, economically grounded realized factor variances. By combining a log-linear specification with adaptive factor selection, the model remains tractable and scalable for large cross-sections yet flexible enough to capture multi-factor volatility dynamics, effectively unifying autoregressive persistence with a multiplicative factor structure.

To motivate the model structure, we document persistence in the key variance components. As established by Ding, Engle, Li & Zheng (2025) and Herskovic, Kelly, Lustig & Van Nieuwerburgh (2016), both the cross-sectional realized variance (CRV) and

[3] As a robustness check, we also implement the selection rule using the RMSE loss metric. The resulting forecasts (reported in the online Appendix C) are qualitatively similar, indicating that our adaptive factor choice is not sensitive to the specific loss function employed.

Table 1: Summary statistics of first-order autocorrelation for factor-level realized variances. The reported values are the 25% quantile (Q1), the median, and the 75% quantile (Q3) of the autocorrelation across the factors under construction.

	Q1	Median	Q3
$\text{corr}(\text{FRV}_t, \text{FRV}_{t-1})$	0.776	0.815	0.846

the multiplicative residual component exhibit substantial long-memory characteristics, reflecting persistent temporal dependencies in asset-specific and market-level volatility dynamics. We extend this analysis to the realized factor variances (FRV) used in our framework. We investigate whether a similar degree of persistence characterizes the realized variances of the high-frequency asset pricing factors included in our forecasting framework.

For each factor in our dataset, we compute the first-order autocorrelation of daily realized variances. Table 1 reports the summary statistics of these autocorrelations, while Figure 2 presents their empirical distribution. Consistent with the hypothesis of long-memory behavior, we find that factor-level realized variances are highly persistent, with a median autocorrelation of 0.815 and an interquartile range spanning from 0.776 to 0.846. These findings corroborate the presence of a strong autoregressive structure in factor volatilities, and support the inclusion of FRV as a dynamic predictor in volatility forecasting models.

The single-factor framework naturally generalizes to the two or three most informative factor variances for each stock. Given the candidate set $\mathcal{K}$, let $\mathcal{K}^\star_{i,t} = \{k^\star_{i,t,1}, k^\star_{i,t,2}, \ldots, k^\star_{i,t,K^\star}\} \subseteq \mathcal{K}$ denote the ordered set of factors selected for stock i at date t, where $K^\star = \text{card}(\mathcal{K}^\star_{i,t}) \in \{2, 3\}$ is the number of selected factors. The elements are ordered by in-sample QLIKE loss with lowest first so that $k^\star_{i,t,1}$ is the top-ranked factor.

In this multi-factor framework, the forecasting equation (7) becomes

$$\log RV_{i,t+h} = \beta_0 + \sum_{z\in\{d,w,m\}} \left(\beta^{(z)}_{CRV} \log CRV^{(z)}_t + \beta^{(z)}_{\xi} \log \xi^{(z)}_{i,t} + \sum_{k\in\mathcal{K}^\star_{i,t}} \beta^{(z)}_k \log FRV^{(z)}_{k,t} \right) + \varepsilon_{i,t+h}. \tag{9}$$

We restrict $K^\star$ to $\{2, 3\}$ in order to balance explanatory power, estimation stability,

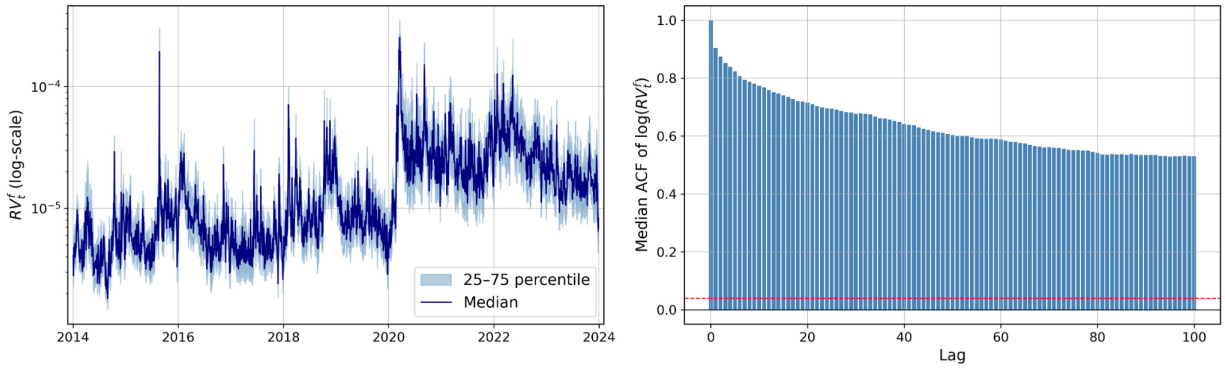

Figure 2: Median persistence of factor realized variances. The left panel plots the daily first-order autocorrelation with its interquartile range (25–75%), while the right panel shows the median autocorrelation function across all high-frequency factors. The red dashed line marks the 5% critical value.

and interpretability; because realized factor variances share considerable information, performance saturates with only a few top factors.[4] The selection routine proceeds iteratively, searching for the next volatility predictor that adds the most explanatory power. First, we scan the full universe of candidate series and retain the factor $k^\star_{i,t,1}$ that delivers the lowest in-sample QLIKE (or RMSE) loss over the rolling selection window S. Holding $k^\star_{i,t,1}$ fixed, we re-estimate the model after adding each remaining candidate one at a time and select the factor $k^\star_{i,t,2}$ that achieves the largest additional loss reduction; this completes the specification with $K^\star = 2$. If a third predictor is allowed, we repeat the exercise once more: evaluate every unused candidate conditional on $\{k^\star_{i,t,1}, k^\star_{i,t,2}\}$, and retain $k^\star_{i,t,3}$, i.e., the factor that yields the largest additional loss reduction. Thus, each additional factor variance is chosen strictly for the marginal improvement it brings.

Forecasts for time $t+h$ are formed at the close of day t using CRV_t, $\xi_{i,t}$, and $FRV_{k,t}$, $k \in \mathcal{K}^\star_{i,t}$, together with their weekly and monthly aggregates. All quantities are $\mathcal{F}_t$-measurable, so the procedure uses only information available by date t and does not use those beyond t.

[4]We also estimated extensions with $K^\star \geq 4$. In our forecasting exercises these higher-dimensional specifications did not yield additional gains in out-of-sample accuracy and, in most cases, delivered weaker performance than the corresponding $K^\star = 3$ model. To preserve conciseness, we do not report these results.

2.3 Model Validity I: Adaptive Selection

We now establish the theoretical foundations of the log-HAR framework with adaptive factor volatilities introduced in Subsections 2.1 and 2.2. We prove in Theorem 1 that the QLIKE-based selector consistently identifies the true predictive factor, verifying the asymptotic correctness of the selection rule. We also show how the selection reliability is affected by the variance of measurement error proxied by realized quarticity in Theorem 2. The results motivate and support the validity of our use of robust realized variance estimators and the factor-volatility block.

We suppose the true data generating process (DGP) for the log integrated variance is

$$\log(IV_{i,t}) = \beta_0 + \beta^{CRV} \log(CIV_t) + \beta^{\xi} \log(\xi_{i,t}^{IV}) + \sum_{k \in \mathcal{K}} \gamma_{k,t} \cdot \log(FIV_{k,t}) + \varepsilon_{i,t}, \tag{10}$$

with $CIV_t := N^{-1} \sum_i IV_{i,t}$, $\xi_{i,t}^{IV} := IV_{i,t}/CIV_t$, and

$$FIV_{k,t} := \int_{t-1}^{t} \sigma_{k,s}^2 \, ds, \tag{11}$$

where $\sigma_{k,s}^2$ is the instantaneous variance of the factor's return at time s and $\gamma_{k,t}$ is the loading on $\log(FIV_{k,t})$, nonzero only for the unique active factor on the window, see Assumption A below. The observed quantities $\log(CRV_t)$, $\log(\xi_{i,t})$ and $\log(FRV_{k,t})$ are proxies for these latent terms. If factor k is a tradable linear portfolio with predictable weights $w_{k,s}$ on constituents with instantaneous covariance Σ_s, then $\sigma_{k,s}^2 = w_{k,s}' \Sigma_s w_{k,s}$ and $FIV_{k,t} = \int_{t-1}^{t} w_{k,s}' \Sigma_s w_{k,s} \, ds$. This justifies the interpretation of $FRV_{k,t}$ as a high-frequency proxy for the latent quantity in the data generating process.

Notations and Preliminaries. We introduce some notations we shall use throughout. Fix a stock i, a forecast origin t, and a forecasting horizon h, and recall that $\mathcal{K}$ is the index set for the candidate factors. In our empirical application $K = 287$, where $K = \text{card}(\mathcal{K}) < \infty$. For a candidate factor $k \in \mathcal{K}$, we define the parameter vector of the corresponding model as $\theta := (\beta_0, \beta_{CRV}, \beta_{\xi}, \beta_k) \in \Theta_k$ and the vector of predictors available at time s as $X_{i,s}^k := \Big(\log(CRV_s), \log(\xi_{i,s}), \log(FRV_{k,s})\Big)$, augmented with their weekly and monthly

aggregates. We denote by $m(X_{i,s}^k;\theta)$ the log-linear predictor for stock i including factor k, so that

$$\log RV_{i,s+h} \;=\; m(X_{i,s}^k;\theta) \;+\; \varepsilon_{i,s+h}, \qquad s \in \{t-S,\dots,t-1\}, \tag{12}$$

cf. (7). Let the level target and its implied forecast be, respectively, $V_{i,s} := RV_{i,s+h} > 0$ and $\widehat{V}_{i,s}^k(\theta) := \exp\{\, m(X_{i,s}^k;\theta)\,\} > 0$. We refer to the single-observation QLIKE loss using

$$\ell\Big(V_{i,s}, X_{i,s}^k;\theta\Big) \;:=\; \log\left(\frac{\widehat{V}_{i,s}^k(\theta)}{V_{i,s}}\right) + \frac{V_{i,s}}{\widehat{V}_{i,s}^k(\theta)} - 1, \tag{13}$$

which is well-defined since $V_{i,s} > 0$ and $\widehat{V}_{i,s}^k(\theta) > 0$, and write the sample loss

$$L_{i,t}(k;\theta) \;=\; \frac{1}{S}\sum_{s=t-S}^{t-1} \ell\Big(V_{i,s}, X_{i,s}^k;\theta\Big)\,. \tag{14}$$

Finally, let

$$\mathcal{L}_{i,t}(k) \;:=\; \inf_{\theta\in\Theta_k} \mathbb{E}\Big[\, \ell\Big(V_{i,s}, X_{i,s}^k;\theta\Big) \;\Big|\; \mathcal{F}_{t-1}\Big]$$

denote the minimized (conditional) population loss for factor $k \in \mathcal{K}$.

We impose the following regularity conditions:

Assumption A.

A1. *For each stock i and time t, there is a single true active factor, $k_{i,t}^\star$, which remains unchanged over $\{t-S,\dots,t-1\}$. That is, the set of true non-zero coefficients $\mathcal{S}_{i,t} = \{k \in \mathcal{K}; \gamma_{k,t} \neq 0\}$ has only one element so that card$(\mathcal{S}_{i,t}) = 1$.*

A2. *$\gamma_{k_{i,s}^\star}$ is uniformly bounded away from zero by some positive constant C*

$$\inf_{s\in[t-S,t-1]} |\gamma_{k_{i,s}^\star}| \geq C > 0. \tag{15}$$

A3. *Let $\tilde{Z}_{i,s} = (1, \log(CRV_s), \log(\xi_{i,s}))$. Let $u_{k_{i,s}^\star}$ and $u_{k,s}$ be the residuals from the population regressions of $\log(FRV_{k_{i,s}^\star})$ and $\log(FRV_{k,s})$ on $\tilde{Z}_{i,s}$, respectively. Then, for some $\kappa \in [0,1)$,*

$$\max_{k\in\mathcal{K}\setminus\{k_{i,s}^\star\}} \Big|\mathrm{corr}(u_{k_{i,s}^\star}, u_{k,s})\Big| \leq \kappa < 1. \tag{16}$$

A4. *(i) $\{(V_{i,s}, X_{i,s}^k)\}$ is strictly stationary and α-mixing with $\alpha(\ell) = O(\ell^{-c})$ for some $c > 2$; (ii) for each $k \in \mathcal{K}$, Θ_k is compact and $\theta \mapsto \ell(V_{i,s}, X_{i,s}^k, \theta)$ is measurable and continuous*

a.s.; (iii) there exists an envelope $M_{i,s}$ with $|\ell(V_{i,s}, X_{i,s}^k, \theta)| \le M_{i,s}$ and $\mathbb{E}[M_{i,s}] < \infty$.

Remark. Assumption A1 imposes a unique, locally stable volatility driver within the selection window, which is the minimal identification content needed for a single-factor selector. Assumption A2 is a standard signal-strength condition ensuring that the contribution of the true factor does not vanish on the window. Assumption A3 is a restricted non-collinearity requirement: after partialling out the common block $\tilde{Z}_{i,s} = (1, \log(CRV_s), \log(\xi_{i,s}))$, the innovation in the true factor cannot be replicated by any inactive factor; this rules out near-collinearity of the factor-specific signals once the common predictors are controlled for. Under the log-DGP in (10) and the QLIKE loss $\ell(\cdot)$, it follows that A2 and A3 jointly imply strict separation of the population risks: there exists $\Delta_{\min} > 0$ such that

$$\mathcal{L}_{i,t}(k) - \mathcal{L}_{i,t}(k_{i,t}^\star) \ge \Delta_{\min} \quad \text{for all} \quad k \in \mathcal{K} \setminus \{k_{i,t}^\star\}. \tag{17}$$

Assumption A4 provides a uniform law of large numbers over $k \in \mathcal{K}$ and $\theta \in \Theta_k$: stationarity/mixing, compact parameter spaces, and an integrable envelope imply

$$\max_{k \in \mathcal{K}} \sup_{\theta \in \Theta_k} \left| L_{i,t}(k;\theta) - \mathbb{E}\left[\ell(V_{i,s}, X_{i,s}^k, \theta)\right] \right| \xrightarrow{p} 0, \tag{18}$$

so the sample criteria deviate uniformly little from their population counterparts, see Andrews (1987), Davidson (1994). Consequently, with K fixed and finite, the sample argmin equals the population argmin with probability tending to one, which is the main content of Theorem 1 below.

The strong consistency of our selection procedure is now formally presented. The theorem proves that our QLIKE-based selection rule asymptotically identifies the factor that truly drives volatility. With sufficiently long windows, the probability that the procedure picks the true factor converges to one, as desired.

Theorem 1. *Let $\mathcal{K}$ be the finite set of candidate factors. Suppose $\widehat{k}_{i,t}^\star$ is chosen at time t for stock i according to the minimum QLIKE loss criterion for the model with factor k.*

That is, for each $k \in \mathcal{K}$ and $\theta \in \Theta_k$, define

$$L_{i,t}(k;\theta) := \frac{1}{S}\sum_{s=t-S}^{t-1} \ell(V_{i,s}, X_{i,s}^{k}, \theta) \qquad \textit{and} \qquad \hat{\theta}_k \in \arg\min_{\theta\in\Theta_k} L_{i,t}(k;\theta), \tag{19}$$

and let

$$\widehat{k}_{i,t}^{\star} \in \operatorname*{argmin}_{k\in\mathcal{K}} L_{i,t}(k;\hat{\theta}_k). \tag{20}$$

If Assumptions $A1 - A4$ hold, then the selection procedure identifies the true active factor with probability tending to one, i.e., as $S \to \infty$ with $S = o(T)$,

$$\mathbb{P}\Big(\widehat{k}_{i,t}^{\star} = k_{i,t}^{\star}\Big) \;\to\; 1, \tag{21}$$

where $k_{i,t}^{\star}$ denotes the unique active factor of (10) in the window.

Proof. See online Appendix A. □

2.4 Extension to Multi-factor forward selection

While Theorem 1 establishes consistency for selecting a *single* volatility driver, our empirical procedure selects $K^{\star} \in \{2, 3\}$ factors via a forward rule. We now provide a unified argument showing that the implemented multi-factor selector is consistent under a sparse multi-factor DGP. The single-factor case corresponds to $K_0 = 1$.

Fix stock i and time t. Let the (locally) true set of active factors be $\mathcal{S}_{i,t} \subseteq \mathcal{K}$ with $K_0 := \text{card}(\mathcal{S}_{i,t}) \leq K^{\star}$, stable on the selection window $\{t-S, \ldots, t-1\}$. For any subset $J \subseteq \mathcal{K}$, let $X_{i,s}^{J}$ denote the predictor vector that augments the baseline block $(\log CRV_s, \log \xi_{i,s})$ with the factor-volatility blocks $\{\log FRV_{k,s}\}_{k\in J}$, and define the sample QLIKE criterion

$$L_{i,t}(J;\theta) := \frac{1}{S}\sum_{s=t-S}^{t-1} \ell\Big(V_{i,s}, X_{i,s}^{J}; \theta\Big), \qquad \widehat{L}_{i,t}(J) := \inf_{\theta\in\Theta_J} L_{i,t}(J;\theta), \tag{22}$$

with the minimized population loss

$$\mathcal{L}_{i,t}(J) := \inf_{\theta\in\Theta_J} \mathbb{E}\Big[\ell\Big(V_{i,s}, X_{i,s}^{J}; \theta\Big) \,\Big|\, \mathcal{F}_{t-1}\Big]. \tag{23}$$

The forward QLIKE selector used in the empirical section is the recursion $\widehat{J}_{i,t,0} = \emptyset$ and, for

$m = 1, \ldots, K^\star$,

$$\widehat{k}_{i,t,m} \in \underset{k \in \mathcal{K} \setminus \widehat{J}_{i,t,m-1}}{\arg\min} \; \widehat{L}_{i,t}\Big(\widehat{J}_{i,t,m-1} \cup \{k\}\Big), \qquad \widehat{J}_{i,t,m} := \widehat{J}_{i,t,m-1} \cup \{\widehat{k}_{i,t,m}\}. \tag{24}$$

Assumption A′.

A1′. *For each stock i and time t there exists a set of active factors $\mathcal{S}_{i,t} \subseteq \mathcal{K}$ with $K_0 = \mathrm{card}(\mathcal{S}_{i,t}) \leq K^\star$ that is unchanged over the selection window $\{t-S, \ldots, t-1\}$.*

A2′. *The nonzero loadings on the active factors are uniformly bounded away from zero on the window: $\inf_{s \in [t-S, t-1]} \min_{k \in \mathcal{S}_{i,t}} |\gamma_{k,s}| \geq C > 0$.*

A3′. *There exists $\Delta_{\min} > 0$ such that for any subset $J \subset \mathcal{S}_{i,t}$ with $\mathrm{card}(J) < K_0$,*

$$\min_{k \in \mathcal{K} \setminus \mathcal{S}_{i,t}} \mathcal{L}_{i,t}(J \cup \{k\}) \; - \; \min_{j \in \mathcal{S}_{i,t} \setminus J} \mathcal{L}_{i,t}(J \cup \{j\}) \; \geq \; \Delta_{\min}, \tag{25}$$

which is the multi-factor analogue of what was used for the single-factor selector.

A4′. *Assumption A4 holds uniformly for all model classes indexed by J with $\mathrm{card}(J) \leq K^\star$.*

Proposition 1. *Under Assumption A′ and with finite $\mathcal{K}$, as $S \to \infty$ with $S = o(T)$,*

$$\mathbb{P}\left(\widehat{J}_{i,t,K_0} = \mathcal{S}_{i,t}\right) \to 1. \tag{26}$$

In particular, if $K^\star \geq K_0$, then $\mathbb{P}(\mathcal{S}_{i,t} \subseteq \widehat{J}_{i,t,K^\star}) \to 1$, and if $K^\star = K_0$ then $\mathbb{P}(\widehat{J}_{i,t,K^\star} = \mathcal{S}_{i,t}) \to 1$.

Proof. See online Appendix A. □

Remark. When $K_0 = 1$, the forward selector (24) reduces to single-factor selection and the above result nests Theorem 1.

2.5 Model Validity II: Adaptive Selection under Measurement Error

We now examine how *measurement error in the factor-volatility block* affects the adaptive selection discussed previously. In our empirical implementation, realized factor variance $FRV_{k,t}$ is a noisy proxy for the latent factor integrated variance $FIV_{k,t}$, and the magnitude of this noise is time varying and captured by realized quarticity (FRQ). Theorem 2 below shows that measurement error *shrinks* the population loss gap that drives the selection: the

higher the (conditional) variance of the measurement error for the *true* factor, the lower the probability that the QLIKE-based selector picks it. This provides a formal rationale for using noise-robust high-frequency estimators: by reducing the measurement-error variance $\sigma^2_{k,s}$ in Assumption B1, they increase selection reliability, as formalized in Theorem 2 below.

We impose the following conditions:

Assumption B.

B1. *For each factor $k \in \mathcal{K}$ and each s,*

$$\log(\widehat{FRV}_{k,s}) = \log(FIV_{k,s}) + e_{k,s},$$

where $\widehat{FRV}_{k,s}$ denotes the realized variance estimator used in practice (e.g., RV, realized kernel, pre-averaging etc.). The error $e_{k,s}$ collects discretization error and microstructure noise, and satisfies $\mathbb{E}[e_{k,s} \mid \mathcal{F}_{s-1}] = 0$, $\mathrm{Var}(e_{k,s} \mid \mathcal{F}_{s-1}) = \sigma^2_{k,s}$, and $\mathbb{E}[|e_{k,s}|^{4+\delta}] < \infty$ for some $\delta > 0$. Moreover, $e_{k,s}$ is conditionally uncorrelated with $\tilde{Z}_{i,s} = (1, \log(CRV_s), \log(\xi_{i,s}))$ and with $\log(FIV_{k',s})$ for $k' \neq k$.

B2. *There exist positive constants $0 < c_- \leq c_+ < \infty$ such that*

$$c_- \, FIQ_{k,s} \;\leq\; \sigma^2_{k,s} \;\leq\; c_+ \, FIQ_{k,s},$$

where $FIQ_{k,s}$ is the factor integrated quarticity, for which $FRQ_{k,s}$ is consistent in the sense that $\sup_k \left| FRQ_{k,s}/FIQ_{k,s} - 1 \right| = o_p(1)$ uniformly in k.

B3. *Let $\bar{\sigma}^2_{k^\star_{i,t}} := S^{-1} \sum_{s=t-S}^{t-1} \sigma^2_{k^\star_{i,s}}$. There exists $\bar{\sigma}^2_0 > 0$ such that $\bar{\sigma}^2_{k^\star_{i,t}} \leq \bar{\sigma}^2_0$, and $\bar{\sigma}^2_0$ can be made arbitrarily small.*

B4. *For each $k \in \mathcal{K}$, the population risk $R_{i,t}(k, \theta) := \mathbb{E}[\ell(V_{i,s}, X^k_{i,s}, \theta)]$ is twice continuously differentiable in the linear predictor $m(X^k_{i,s}; \theta)$ in a neighborhood of the optimum, with a uniform lower curvature bound: there exists $c_0 > 0$ such that for all k,*

$$R_{i,t}(k, \theta) - \inf_{\vartheta \in \Theta_k} R_{i,t}(k, \vartheta) \;\geq\; c_0 \, \mathbb{E}\Big[\Big(m(X^k_{i,s}; \theta) - m(X^k_{i,s}; \theta^\circ_k) \Big)^2 \Big],$$

where $\theta^\circ_k \in \arg\min_{\vartheta \in \Theta_k} R_{i,t}(k, \vartheta)$ under latent inputs.

Remark. The assumptions are mild and specify a weak set of conditions under which the theory is valid. Assumption B1 specifies an errors-in-variable structure, and B2 ties the conditional variance of the error to quarticity. Assumption B3 bounds the window-average noise variance and B4 gives strong convexity of the population QLIKE risk in the log predictor.

We have the following result:

Theorem 2. *Under Assumptions A1-A4 and B1-B4, for*

$$\bar{\sigma}^2_{k^\star_{i,t}} := \frac{1}{S}\sum_{s=t-S}^{t-1} \sigma^2_{k^\star_{i,s}} \qquad \text{and} \qquad H_{i,t}(S) := \max_{k\in\mathcal{K}} \sup_{\theta\in\Theta_k} \Big| L_{i,t}(k;\theta) - R_{i,t}(k,\theta) \Big|,$$

there exist a constant independent of S, denoted $C_1 > 0$, such that the following holds: for any incorrect $k \neq k^\star_{i,t}$,

$$R^\eta_{i,t}(k) - R^\eta_{i,t}(k^\star_{i,t}) \;\geq\; \Delta_{\min} - C_1\, \bar{\sigma}^2_{k^\star_{i,t}} + o(\bar{\sigma}^2_{k^\star_{i,t}}), \tag{27}$$

where $\Delta_{\min}$ is as in (17) and $R^\eta_{i,t}(\cdot)$ denotes the population risk evaluated with the noisy factor inputs. In addition, we have

$$\mathbb{P}\Big(\widehat{k}^\star_{i,t} = k^\star_{i,t}\Big) \;\geq\; \mathbb{P}\Big(H_{i,t}(S) \;<\; \tfrac{1}{2}\big[\Delta_{\min} - C_1\, \bar{\sigma}^2_{k^\star_{i,t}}\big]\Big). \tag{28}$$

In particular, if $\bar{\sigma}^2_{k^\star_{i,t}} < \Delta_{\min}/C_1$, then $R^\eta_{i,t}(k) > R^\eta_{i,t}(k^\star_{i,t})$ for all $k \neq k^\star_{i,t}$, and as $S \to \infty$ with $S = o(T)$, we have

$$\mathbb{P}\Big(\widehat{k}^\star_{i,t} = k^\star_{i,t}\Big) \;\to\; 1. \tag{29}$$

Proof. See online Appendix A. □

3 Data

We assemble a comprehensive high-frequency dataset of U.S. equities to support the estimation of realized variance under infill asymptotics. The sample contains all CRSP common shares (share codes 10 or 11) listed on the NYSE, NASDAQ or AMEX (main exchange codes 1, 2 or 3) from January 2, 2014 to December 29, 2023, covering $T = 2516$ trading days and $N = 5370$ unique securities.

Intraday trades and quotes are sourced from the NYSE TAQ database via WRDS and restricted to regular trading hours (09:30–16:00 Eastern Time).[5] Following standard practice, we apply the filters of Barndorff-Nielsen, Hansen, Lunde & Shephard (2009) to remove observations outside regular hours, zero or negative prices, obvious quote errors, and extreme outliers. Additional safeguards remove FINRA Alternative Display Facility prints (exchange "D") and delete trades priced beyond the daily CRSP ask-high or bid-low. Prices are then sampled on a uniform one-second grid between 09:30:00 and 16:00:00 ET using the previous-tick method of Gençay, Dacorogna, Muller, Pictet & Olsen (2001). The resulting high-frequency returns can naturally be aggregated to any lower frequency, allowing a straightforward implementation for any realized variance estimator.

Daily opens, closes, shares outstanding, and delisting returns are obtained from CRSP. We match TAQ symbols to CRSP identifiers (PERMNOs) via the TAQ–CRSP Link table (covering about 98% of the universe).[6] We prioritize CRSP entries: each CRSP record is retained even if no intraday data exist that day, while unmatched TAQ observations are ignored.[7] To ensure consistency around corporate events, we overwrite the 09:30 and 16:00 TAQ prices with the CRSP open and close, and we adjust the closing return when delisting returns occur, in line with Hou, Xue & Zhang (2018).

3.1 High-frequency factors

Leveraging intraday stock information, we replicate a universe of $K = 287$ high-frequency factors. The first block contains the six canonical Market (MKT), Size (SMB), Value (HML), Profitability (RMW), Investment (CMA), and Momentum (UMD) factors, replicated to closely match the definitions in Fama & French (2018). The second block spans 281

[5] We use the SAS code from Holden & Jacobsen (2014) to extract tick-by-tick transactions matched with contemporaneous bid–ask quotes from daily TAQ. Timestamps are recorded in milliseconds until mid-2015 and microseconds thereafter.

[6] Residual cases are matched by eight-digit CUSIP from the TAQ Master file.

[7] Unmatched observations involve fewer than 1% of stocks, almost all nano-caps, so they have little effect on weighted portfolio returns.

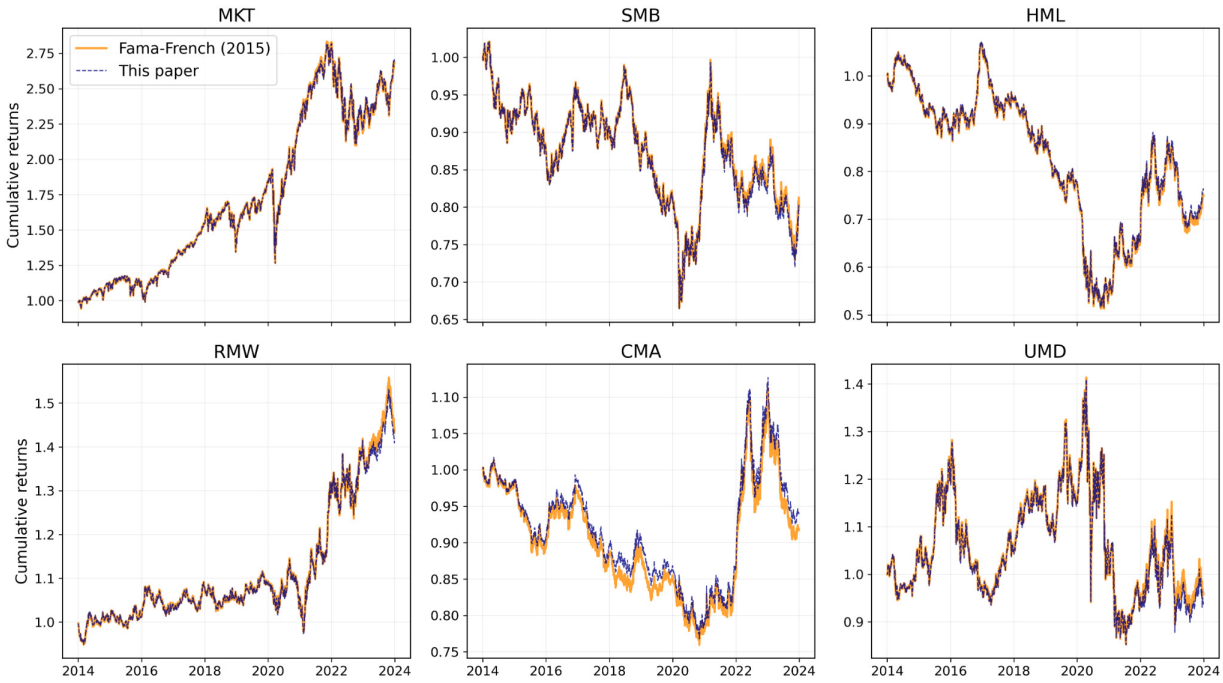

Figure 3: Cumulative gross daily returns of the Fama–French six factors using official data and our high-frequency replication. The orange line corresponds to daily returns sourced from the Kenneth R. French Data Library, while the blue dashed line represents our version constructed from 1-second returns aggregated to daily frequency.

characteristic-sorted portfolios drawn from the large collections of Jensen, Kelly & Pedersen (2023) (JKP) and Chen & Zimmermann (2022) (CZ).[8]

To replicate the Fama–French factors at high frequency, we follow the standard definitions with NYSE breakpoints, double-sorting on size and annual rebalancing in June. Adapting the procedure of Aït-Sahalia, Kalnina & Xiu (2020), we update value-weighted stock returns at 1-second frequency, and compute portfolio returns. Figure 3 shows that our daily aggregation of the replicated factors is virtually indistinguishable from the official data.

The remaining portfolios are constructed from a large cross-section of firm characteristics.[9] At the end of each month, eligible stocks are sorted into terciles by the given characteristic. Following the empirical design in Jensen, Kelly & Pedersen (2023), we

[8]As the Fama–French factors are widely used as benchmarks, we replicate them as faithfully as possible. By contrast, the broader zoo of factors is built under a uniform methodology to ensure comparability across signals.

[9]To avoid duplication where JKP and CZ provide conceptually similar characteristics, we retain the JKP implementation whenever the two series are empirically equivalent (return correlation $> 95\%$) and drop signals with missing data, yielding 153 JKP and 128 CZ unique factors.

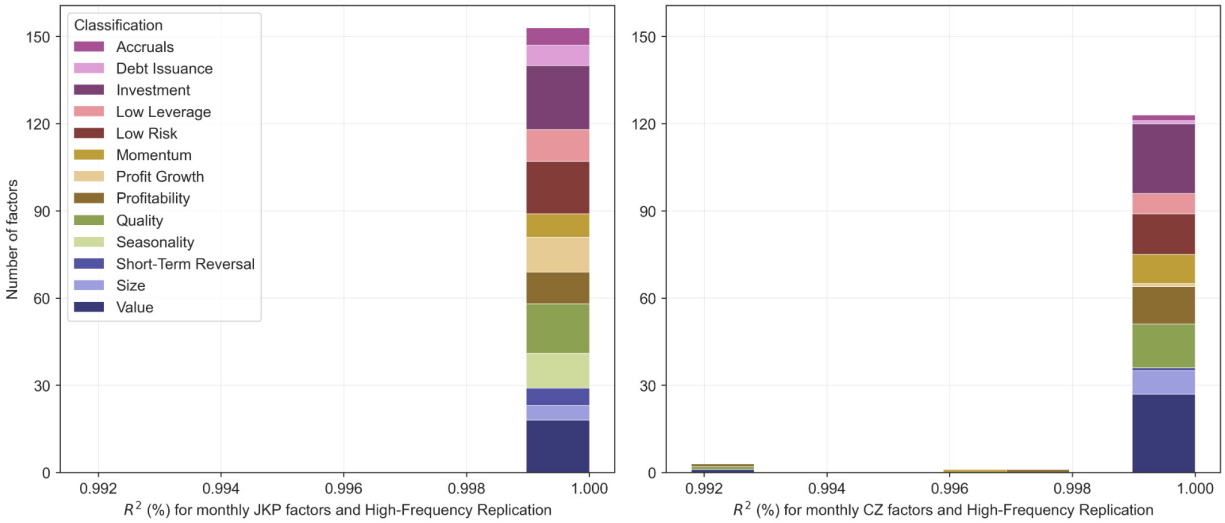

Figure 4: Histogram of R^2 from factor-level monthly return regressions. For each JKP (left) and selected CZ (right) factor, we regress our aggregated high-frequency portfolio returns on the original low-frequency portfolio over 2014–2023. Results are reported by economic cluster using the Jensen, Kelly & Pedersen (2023) taxonomy: CZ factors are mapped to clusters by the highest average correlation between their CAPM-residual returns and those of the JKP factors within each cluster. Higher R^2 indicates closer replication fidelity.

compute value-weighted returns for the top and bottom terciles and form a zero-investment high-minus-low spread held over the subsequent month. To validate our procedure, we compare our replication against the original low-frequency versions by examining the explanatory power of monthly return regressions. Figure 4 summarizes the comparison across JKP and CZ factors, demonstrating high fidelity for the whole replication.

A compendium of the sample factors and more details on the dataset are reported in the online Appendix B. To ensure comparability with the literature and computational tractability, we compute CRV_t and related quantities using the full-universe definitions as in Ding, Engle, Li & Zheng (2025). As the factors and CRV_t are constructed from the same stock universe as the dependent variable, stock i mechanically enters both CRV_t and the factor realized variance ($FRV_{k,t}$) used to forecast $RV_{i,t+h}$. This own-observation inclusion has negligible impact on forecasts: the leave-one-out measure $CRV_t^{(-i)} = (N-1)^{-1}\sum_{j\neq i} RV_{j,t}$ differs from CRV_t by at most $O(1/N)$, and the factor portfolios are well diversified. Accordingly, we treat the effect as negligible.

4 Empirical evidence

The primary interest of the empirical analysis is the out-of-sample forecasting performance of the proposed model. The results of our framework are tested using twelve alternative realized volatility estimators, selected to encompass a broad spectrum of methodological approaches and robustness features. These include the classical realized variance (RV) of Andersen & Bollerslev (1998) at 5-minute and 1-minute frequency, and its 5-minute subsampled version. We consider the bipower variation (BPV) introduced by Barndorff-Nielsen & Shephard (2004) and its staggered definition (Andersen et al. 2007), the realized kernel developed by Barndorff-Nielsen, Hansen, Lunde & Shephard (2008), the truncated realized variance (TRV) of Mancini (2009) and the pre-averaged measures (PRV, PBV) of Christensen, Oomen & Podolskij (2014). We also include more recent contributions like the differenced-return variance (DV) of Andersen, Li, Todorov & Zhou (2023) and the nonparametric price duration variance (NPDV) proposed by Hong, Nolte, Taylor & Zhao (2023).

While all results are evaluated across the complete set of estimators to ensure robustness, we adopt the candlestick variance (or wick variance, WV) by Li, Nolte, Nolte & Yu (2025) as the benchmark throughout the analysis due to its desirable properties. Specifically, the WV estimator provides resilience against microstructure noise and accommodates jumps and extreme price movements, remaining unbiased and robust under a wide range of market conditions. These features make it especially well-suited for high-frequency volatility estimation and motivate its primary role in our results.

Our forecasts are constructed using a recursive estimation approach. For each stock i and day t, we estimate the coefficients of the model in 9 by ordinary least squares over a rolling window of length $L = 1260$. We generate h-day-ahead forecasts for the log-realized variance, $\widehat{\log RV_{i,t+h}}$, using the estimated parameters and the daily, weekly, and monthly lags of the model explanatory components: the market-wide realized variance (CRV_t), the

corresponding idiosyncratic component ($\xi_{i,t}$), and the log realized variance of the factors ($FRV_{k^\star_{i,t}}$) chosen for the best performance over the selection window $S = 252$.

We assess forecasts across three different forecasting horizons $h \in \{1, 5, 22\}$, corresponding to predictions for one trading day, week and month ahead. Forecast accuracy is assessed using the quasi-likelihood (QLIKE) loss function, defined by Patton (2011) and formulated by Bollerslev, Patton & Quaedvlieg (2016) as

$$QLIKE^i = \frac{1}{Q}\sum_{q=1}^{Q}\left(\log\frac{\widehat{RV}_q^i}{RV_q^i} + \frac{RV_q^i}{\widehat{RV}_q^i} - 1\right), \tag{30}$$

where $\widehat{RV}_q^i$ and RV_q^i denote the forecasted and realized variances of stock i at horizon q, and Q is the number of out-of-sample observations.

4.1 Forecasting performance

We evaluate the forecast accuracy over the sample of $N = 1041$ CRSP stocks having less than 20% zero 5-minute returns between 2014 and 2023, with an evaluation period of $Q = 1213$ trading days. Realized variance is measured using the candlestick estimator of Li et al. (2025). Our factor-augmented specifications in (9) are denoted by $HAR\text{-}1F$, $HAR\text{-}2F$, and $HAR\text{-}3F$, corresponding respectively to $\mathcal{K}^\star = 1, 2, 3$. We compare their forecasting performance against widely adopted benchmarks, namely the standard HAR by Corsi (2009), the quarticity-adjusted HARQ of Bollerslev, Patton & Quaedvlieg (2016) and the asymmetric SHAR (Patton & Sheppard 2015). We also include the market augmented HAR of Hizmeri, Izzeldin, Nolte & Pappas (2022) and the multiplicative volatility by Ding, Engle, Li & Zheng (2025) with common RV (MFV-CRV) and first principal component (MFV-PC1).[10]

In Table 2, the models we propose attain the lowest QLIKE across every quantile of the cross-section. The performance strengthens with the number of selected factors and with

[10]To isolate the incremental contribution of realized factor variances, we also consider the restricted version of our forecasting equation obtained by retaining only the CRV and ξ HAR components. In our out-of-sample evaluation, this baseline delivers forecast accuracy that is closely aligned with the MFV benchmarks. We omit it from the main results to avoid redundancy.

Table 2: Cross-sectional QLIKE loss distributions of different forecasting models and horizons, using the candlestick variance estimator of Li et al. (2025). Entries report the 25th percentile (Q1), mean, median (50th), and 75th percentile across $N = 1041$ stocks over an evaluation period of $Q = 1213$ trading days. In every row, bold entries indicate the lowest value.

	Forecasting models								
	HAR	SHAR	HARQ	HAR-MKT	MFV-CRV	MFV-PC1	HAR-1F	HAR-2F	HAR-3F
Panel A: $h = 1$									
Q1	0.1406	0.1414	0.1372	0.1415	0.1359	0.1340	0.1293	**0.1288**	**0.1288**
Median	0.1701	0.1738	0.1668	0.1728	0.1650	0.1632	0.1570	**0.1559**	0.1565
Mean	0.1816	0.1867	0.1779	0.1844	0.1767	0.1738	0.1674	0.1666	**0.1665**
Q3	0.2142	0.2218	0.2102	0.2175	0.2077	0.2041	0.1961	**0.1946**	0.1949
Panel B: $h = 5$									
Q1	0.2335	0.2427	0.2328	0.2344	0.2286	0.2228	0.2003	0.1950	**0.1904**
Median	0.2740	0.2827	0.2719	0.2747	0.2630	0.2563	0.2325	0.2246	**0.2202**
Mean	0.2900	0.2987	0.2865	0.2897	0.2768	0.2715	0.2455	0.2373	**0.2313**
Q3	0.3270	0.3363	0.3230	0.3269	0.3074	0.2990	0.2746	0.2656	**0.2573**
Panel C: $h = 22$									
Q1	0.4920	0.4988	0.4927	0.4877	0.4841	0.4860	0.3434	0.3096	**0.2654**
Median	0.5941	0.5967	0.5928	0.5834	0.5841	0.5797	0.3972	0.3527	**0.3039**
Mean	0.6179	0.6218	0.6165	0.6105	0.6065	0.6079	0.4133	0.3644	**0.3145**
Q3	0.7061	0.7163	0.7033	0.6997	0.6948	0.6993	0.4627	0.4094	**0.3498**

the length of the forecasting horizon, with the three-factor model lowering the average loss metric by up to 48% relative to the best alternative. Figure 5 complements these results by displaying the full distributions. The factor specifications exhibit a clear downward shift and tighter interquartile ranges relative to all benchmarks, with compressed upper tails that indicate robustness to outliers. The improvements are monotone, confirming that the inclusion of multiple volatility drivers yield lower and more stable forecasting loss. Our findings remain qualitatively unchanged across all alternative estimators considered[11].

We next evaluate the model performance at the individual stock level and assess the statistical relevance of the observed forecast improvements. Specifically, we compare the QLIKE loss of the HAR-1F model against that of each benchmark across the stock universe[12].

We quantify the fraction of stocks for which our factor model yields a lower QLIKE as

[11]The online Appendix C verifies robustness to the choice of realized volatility estimator, forecasting horizon $h \in \{1, 5, 22\}$, and scoring rule.

[12]Figure 5 indicates that the multi-factor specifications (HAR-2F, HAR-3F) achieve further reductions in QLIKE relative to HAR-1F and the benchmarks: the one-factor outputs therefore provide a conservative result.

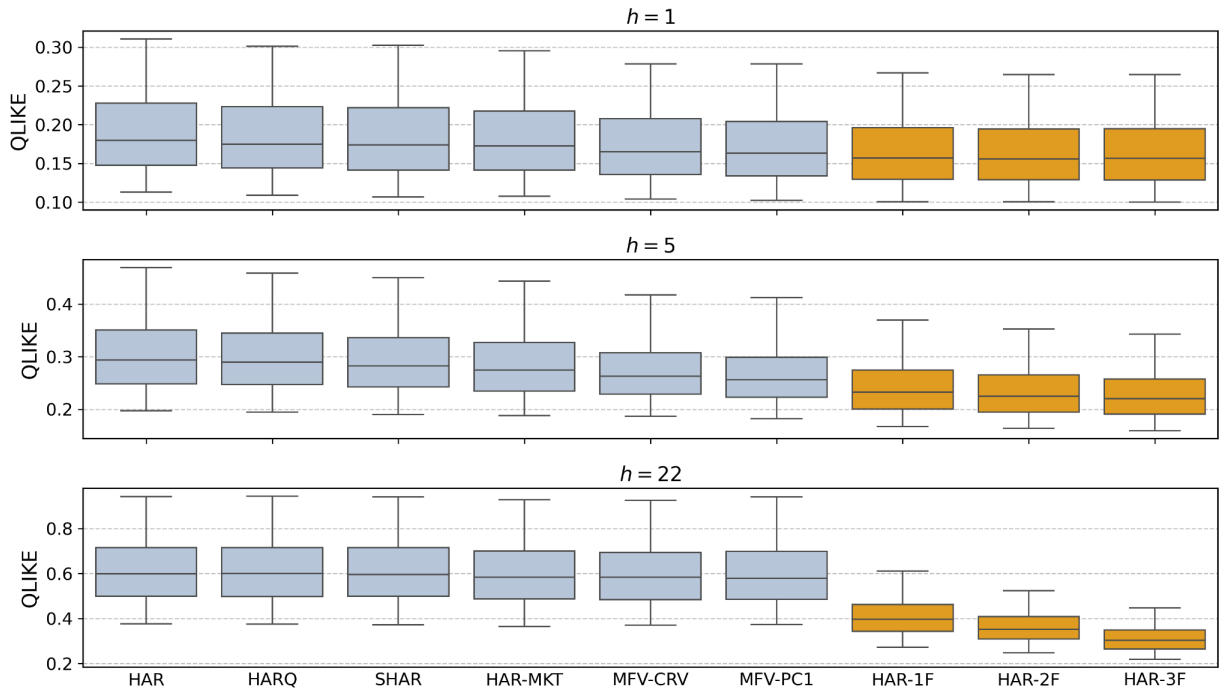

Figure 5: Distribution of QLIKE losses across models and forecast windows. Each boxplot summarizes the QLIKE values across all stocks using the WV estimator (Li et al. 2025). The models include traditional benchmarks (HAR, HARQ, SHAR), the market-HAR, the MFV model with common RV and PC1, the proposed one-factor model and its multi-factor extensions. Boxes represent the 25th, 50th, and 75th percentiles, while whiskers extend to the 1st and 99th percentiles.

the raw outperformance

$$\text{Out.perf.} = \frac{1}{N}\sum_{i=1}^{N} \mathbb{1}\left\{QLIKE^{i}_{HAR-1F} < QLIKE^{i}_{bm}\right\}, \tag{31}$$

where $QLIKE^{i}_{HAR-1F}$ and $QLIKE^{i}_{bm}$ denote the QLIKE losses of stock i for the proposed model and the benchmark model, respectively. To further examine the statistical significance of the forecast differentials, we implement the Diebold & Mariano (2002) test for each stock. We evaluate the null hypothesis of equal predictive accuracy using a one-sided test on the difference between the QLIKE values of the proposed and benchmark models.[13] Finally, we report the proportion of stocks for which our model delivers statistically significant forecast improvements

[13]The null and alternative hypotheses are $H_0 : \mathbb{E}(QLIKE_{HAR-1F,t} - QLIKE_{bm,t}) \geq 0$ against $H_1 : \mathbb{E}(QLIKE_{HAR-1F,t} - QLIKE_{bm,t}) < 0$. The test statistic is computed as $\bar{d}/\hat{\sigma}(\bar{d})$, where $\bar{d} = \sum_{q=1}^{Q} d_q/Q$ is the average loss differential $d_q = QLIKE_{HAR-1F,q} - QLIKE_{bm,q}$, and $\hat{\sigma}(\bar{q})$ is its heteroskedasticity-autocorrelation consistent (HAC) standard error.

$$\text{Sig.Out.perf.} = \frac{1}{N}\sum_{i=1}^{N} \mathbb{1}\left\{p_i < \alpha\right\}, \tag{32}$$

where p_i is the p-value of the DM test for stock i, and $\alpha = 5\%$ is the significance threshold.

Table 3 reports the proportion of cases for which our factor augmented model statistically outperforms each benchmark. The significative outperformances show pervasive gains for our one-factor specification relative to the traditional benchmarks across any volatility estimator and forecasting horizon. For the weekly horizon, more than 95% of stocks exhibit significantly lower QLIKE using the HAR-1F rather than the traditional HAR model or its univariate variants. At the monthly horizon, the improvement increases to more than 99% in every case. At the daily horizon significance remains high but less pronounced, with differences depending on the selected estimator. The pattern is stable across noise and jump robust measures, which indicates that the results are not driven by the choice of the realized specification. Overall, our model's superior statistical performance is robust across different estimators, benchmarks, and forecasting horizons: these results underscore the added value of flexible factor selection in capturing time-varying heterogeneity in volatility sources.

4.2 Return and volatility factor selections

We examine whether, within a broad factor universe with $K = 287$ characteristics, the same cross-sectional signal tends to drive both expected returns and realized variances for a given stock and day. For each (i, t), we compare the factor selected by an adaptive return model to the volatility driver selected by our volatility framework. The return selector is a one-factor specification with a cross-sectional component and a single characteristic factor return, chosen to minimize the sum of squared errors over the most recent $S = 252$ observations within a rolling window of length $L = 1260$. The volatility selector is the proposed HAR-1F model.

We map each factor to one of thirteen economic clusters using the taxonomy of Jensen,

Table 3: Significative outperformance proportion of the one-factor model (HAR-1F) over the benchmark models across forecasting horizons and volatility estimators between 2014 and 2023. Entries report the percentages of stocks for which the QLIKE loss of the HAR-1F model is significantly lower than the benchmark performance at the 5% significance level of the Diebold–Mariano test.

	Volatility estimators											
Benchmarks	RV5	$RV5_{ss}$	RV1	BPV	BPV_{stag}	RK	TRV	PRV	PBV	NPDV	DV	WV
Panel A: $h = 1$												
HAR	94.7	94.0	92.8	95.8	95.5	97.1	96.0	96.3	96.4	95.6	96.6	96.8
HARQ	64.2	87.9	91.9	96.0	95.9	97.0	95.2	92.5	92.0	94.8	95.5	88.8
SHAR	78.4	93.7	84.1	79.8	82.2	78.0	90.5	73.0	71.2	96.0	84.6	75.2
HAR-MKT	83.7	84.4	85.4	88.0	87.2	89.3	91.3	88.0	90.6	64.6	89.5	90.9
MFV-CRV	64.1	59.6	62.3	65.5	61.5	76.9	57.6	75.6	73.6	57.8	56.4	78.9
MFV-PC1	45.0	92.0	41.8	95.3	95.4	56.2	95.2	96.3	96.9	94.9	96.0	53.1
Panel B: $h = 5$												
HAR	98.2	97.8	97.7	98.1	98.0	97.9	97.6	97.9	97.6	99.1	98.5	99.2
HARQ	97.0	97.8	97.6	98.2	98.2	98.0	97.6	97.4	97.0	99.0	98.0	98.3
SHAR	97.9	97.7	97.0	97.3	97.1	97.3	96.9	96.0	95.1	98.4	97.0	96.7
HAR-MKT	95.1	95.7	96.7	95.3	95.5	95.1	96.2	93.7	93.3	96.6	96.9	97.7
MFV-CRV	84.9	83.0	87.4	85.8	84.8	87.1	86.7	85.3	85.1	93.3	87.9	93.1
MFV-PC1	77.0	92.6	79.0	93.3	93.0	78.2	92.0	94.0	93.7	96.8	93.2	85.1
Panel C: $h = 22$												
HAR	99.1	99.2	99.2	99.5	99.4	99.2	99.8	99.2	99.3	99.7	99.8	99.3
HARQ	99.1	99.2	99.2	99.5	99.5	99.2	99.8	99.2	99.2	99.6	99.8	99.3
SHAR	99.1	99.2	99.2	99.4	99.5	99.2	99.8	99.3	99.3	99.5	99.8	99.3
HAR-MKT	99.3	99.2	99.5	99.5	99.7	99.3	99.8	99.3	99.4	99.6	99.8	99.6
MFV-CRV	99.2	99.4	99.4	99.4	99.6	99.3	99.8	99.4	99.4	99.8	99.6	99.6
MFV-PC1	99.4	99.0	99.5	99.2	99.4	99.4	99.7	98.9	99.1	99.5	99.5	99.6

Kelly & Pedersen (2023). JKP portfolios inherit their original labels, while CZ portfolios are assigned to clusters by the highest average correlation between their CAPM-residual returns and those of the JKP factors within each cluster. Figure 6 displays, for each trading day, the distribution across clusters of the factor selected by the return model (left) and by the HAR-1F volatility model (right), expressed as the share of stocks selecting a factor from each cluster; the stacked areas sum to 100% per day.

To quantify alignment at the stock–day level, let $k_{i,t}^{R}$ denote the factor chosen by the return selector and $k_{i,t}^{V}$ the factor chosen by the volatility selector, and let $c(k) \in \{1, \ldots, 13\}$

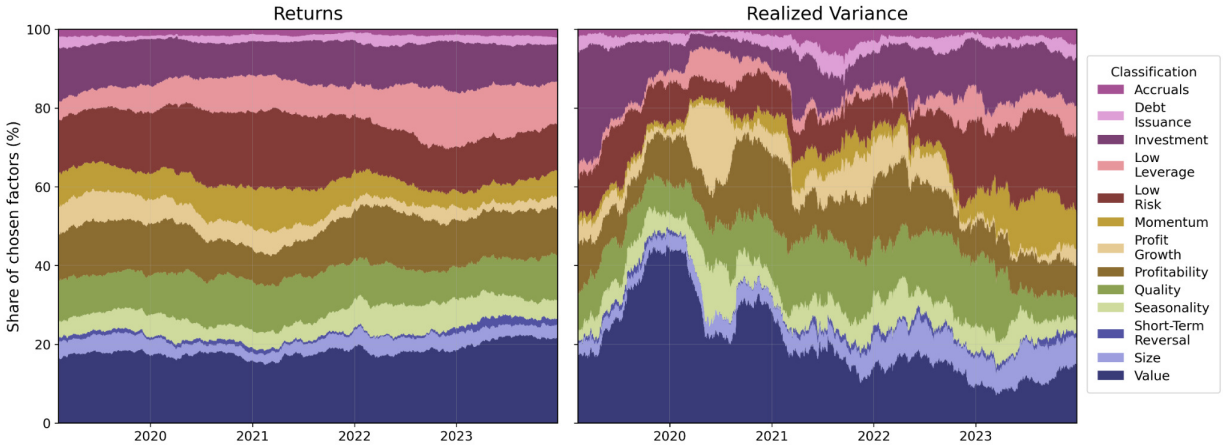

Figure 6: Time series of cluster shares implied by the adaptive selectors. Left: share of stocks whose return model selects a factor from each cluster. Right: corresponding shares for the HAR-1F volatility model using the candlestick estimator of Li et al. (2025). Stacked areas sum to 100% per day over a period of length $Q = 1213$ trading days.

map factors to clusters. Define the daily proportions

$$p_t^{\text{factor}} = \frac{1}{N}\sum_{i=1}^{N} \mathbb{1}\left\{ k_{i,t}^R = k_{i,t}^V \right\} \qquad \textit{and} \qquad p_t^{\text{cluster}} = \frac{1}{N}\sum_{i=1}^{N} \mathbb{1}\left\{ c\left(k_{i,t}^R\right) = c\left(k_{i,t}^V\right) \right\}.$$

Across the sample, the median of p_t^{factor} is below 1%, and the median of p_t^{cluster} is below 10%, indicating that exact factor matches and even same-cluster matches are rare in the cross-section on any given day. Consequently, similarities between the left and right panels of Figure 6 reflect aggregate shifts in cluster composition across stocks, not synchronized selections for the same stocks. These findings are consistent across forecasting horizons. When the volatility factor is selected for weekly and monthly forecasts, the chosen clusters concentrate on a smaller subset, with the strongest concentration during high-volatility regimes. Results are qualitatively robust to alternative volatility estimators and to loss functions beyond QLIKE.

4.3 Economic significance

Beyond statistical accuracy, an important question is whether improvements in volatility forecasting translate into economically meaningful gains. To address this, we adopt the utility-based evaluation framework introduced by Bollerslev, Hood, Huss & Pedersen (2018), which links variance forecasts to the performance of volatility-managed investment strategies.

Table 4: Economic-value comparison between the HAR-1F model and the benchmarks across daily to monthly forecasting horizons, using the WV estimator of Li et al. (2025). For each benchmark, entries report the cross-sectional percentage of stocks for which HAR-1F delivers higher annualized utility than the benchmark (Out.perf.), the percentage for which the improvement is statistically significant at the 5% level using a two-sided Diebold–Mariano test (Sig.Out.perf.), and the percentage for which the utility difference exceeds 1 basis point per year (Diff $\geq$ 1 bp). Values in parentheses give the corresponding percentages in which the benchmark outperforms the HAR-1F model under the same criterion.

Benchmarks	Out.perf. (%)	Sig.Out.perf. (%)	Diff $\geq$ 1 bp (%)
Panel A: $h = 1$			
HAR	99.3 (0.7)	96.8 (0.0)	98.4 (0.5)
HARQ	97.5 (2.5)	88.8 (0.5)	94.1 (1.5)
SHAR	91.9 (8.1)	75.2 (0.8)	88.0 (3.6)
HAR-MKT	99.2 (0.8)	90.9 (0.0)	94.4 (0.5)
MFV-CRV	93.9 (6.1)	78.9 (0.2)	82.7 (0.4)
MFV-PC1	85.0 (15.0)	53.1 (2.3)	68.8 (1.7)
Panel B: $h = 5$			
HAR	99.9 (0.1)	99.2 (0.1)	99.8 (0.1)
HARQ	99.9 (0.1)	98.3 (0.1)	99.8 (0.1)
SHAR	99.3 (0.7)	96.7 (0.1)	99.3 (0.2)
HAR-MKT	99.9 (0.1)	97.7 (0.1)	99.6 (0.1)
MFV-CRV	98.9 (1.1)	93.1 (0.2)	98.0 (0.5)
MFV-PC1	95.9 (4.1)	85.1 (1.2)	94.6 (2.3)
Panel C: $h = 22$			
HAR	99.9 (0.1)	99.3 (0.0)	99.8 (0.1)
HARQ	99.9 (0.1)	99.3 (0.0)	99.8 (0.1)
SHAR	99.8 (0.2)	99.3 (0.0)	99.8 (0.2)
HAR-MKT	99.9 (0.1)	99.6 (0.0)	99.8 (0.2)
MFV-CRV	100.0 (0.0)	99.6 (0.0)	99.9 (0.0)
MFV-PC1	100.0 (0.0)	99.6 (0.0)	99.9 (0.0)

Holding stock i, the investor's utility expressed in annualized percentage terms is

$$U^i = \frac{1}{Q}\sum_{q=1}^{Q}\left(8\%\sqrt{\frac{RV_q^i}{\widehat{RV}_q^i}} - 4\%\frac{RV_q^i}{\widehat{RV}_q^i}\right), \tag{33}$$

where RV_q^i and $\widehat{RV}_q^i$ denote, respectively, the realized variance on day q and its forecasted value. The first term captures the expected return component, inversely scaled by conditional volatility, while the second term penalizes variance under a mean-variance investor framework with a risk aversion level implied by a 4% volatility penalty.

Table 4 summarizes the economic performance of our factor-switching model relative

to the considered benchmarks across all forecasting horizons. For each comparison, we report three measures: (i) the percentage of stocks for which the factor model yields higher utility than the benchmark (Out.perf.), (ii) the percentage of stocks for which the utility improvement is statistically significant at the 5% level based on a DM test (Sig.Out.perf.), and (iii) the percentage of stocks for which the utility difference exceeds 1 basis point annually, a threshold used to capture economically significant improvements.

The results demonstrate that the predictive gains of our model translate into substantial economic value. Across all benchmark comparisons and forecast horizons, the factor-switching model consistently improves investor utility for the large majority of stocks, with improvements that are both statistically significant and economically meaningful. These findings reinforce the practical value of incorporating adaptive factor selection in the modeling of stock volatility, and remind that more accurate volatility forecasts yield tangible benefits in portfolio outcomes.

5 Conclusion

This paper introduces a novel volatility forecasting framework that integrates high-frequency factor information with a dynamic selection mechanism to improve the prediction of individual stock variances. By extending a multiplicative volatility model with heterogeneous autoregressive lags and leveraging a broad cross-section of 287 factor volatilities, our approach achieves substantial gains in forecast accuracy relative to standard benchmarks.

The key contribution lies in demonstrating that volatility is not uniformly driven by a fixed set of risk sources but is instead shaped by time-varying, asset-specific exposures to distinct factor volatilities. In this sense, our model attempts to contribute to the methodological gap between volatility modeling and asset pricing.

Beyond the econometric gains, our results underscore the feasibility and informativeness

of replicating factor portfolios at ultra-high frequencies, opening a path toward more granular assessments of systematic volatility. Future research could extend this framework to multivariate volatility modeling and examine the implications for portfolio risk management and derivative pricing.

Supplemental Materials

In the online Appendix we provide the proofs of the main theoretical results, document additional details on the high-frequency dataset and factor construction, and report supplementary tables and figures for the robustness checks across realized volatility estimators, estimation-window lengths, forecasting horizons, and loss functions.

Data Availability Statement

The data that support the findings of this study are available from Wharton Research Data Services. Restrictions may apply to the availability of these data, which were used under license for this study.

Acknowledgments

We thank Nicola Fusari, Yifan Li, Aleksey Kolokolov, Olga Kolokolova, Manh Pham, Roberto Renò, Shifan Yu, Xinyi Zhang (discussant), as well as participants at the Financial Econometrics Conference in Honour of Stephen Taylor and the 2nd Financial Fraud, Misconduct and Market Manipulation Conference, and seminars at various institutions, for helpful comments and discussions.

Funding

This research received no external funding.

Disclosure statement

The authors report there are no competing interests to declare.

References

Aït-Sahalia, Y. & Jacod, J. (2014), *High Frequency Financial Econometrics*, Princeton University Press.

Aït-Sahalia, Y., Kalnina, I. & Xiu, D. (2020), 'High-frequency factor models and regressions', *Journal of Econometrics* **216**(1), 86–105.

Andersen, T. G. & Bollerslev, T. (1998), 'Answering the skeptics: Yes, standard volatility models do provide accurate forecasts', *International economic review* pp. 885–905.

Andersen, T. G., Bollerslev, T. & Diebold, F. X. (2007), 'Roughing it up: Including jump components in the measurement, modeling, and forecasting of return volatility', *The review of economics and statistics* **89**(4), 701–720.

Andersen, T. G., Li, Y., Todorov, V. & Zhou, B. (2023), 'Volatility measurement with pockets of extreme return persistence', *Journal of Econometrics* **237**(2), 105048.

Andrews, D. W. K. (1987), 'Consistency in nonlinear econometric models: A generic uniform law of large numbers', *Econometrica* **55**(6), 1465–1471.

Asai, M., Caporin, M. & McAleer, M. (2015), 'Forecasting value-at-risk using block structure multivariate stochastic volatility models', *International Review of Economics & Finance* **40**, 40–50.

Atak, A. & Kapetanios, G. (2013), 'A factor approach to realized volatility forecasting in the presence of finite jumps and cross-sectional correlation in pricing errors', *Economics Letters* **120**(2), 224–228.

Barigozzi, M. & Hallin, M. (2017), 'Generalized dynamic factor models and volatilities: estimation and forecasting', *Journal of Econometrics* **201**(2), 307–321.

Barigozzi, M. & Hallin, M. (2020), 'Generalized dynamic factor models and volatilities: Consistency, rates, and prediction intervals', *Journal of Econometrics* **216**(1), 4–34.

Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A. & Shephard, N. (2008), 'Designing realized kernels to measure the ex post variation of equity prices in the presence of noise', *Econometrica* **76**(6), 1481–1536.

Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A. & Shephard, N. (2009), 'Realized kernels in practice: Trades and quotes'.

Barndorff-Nielsen, O. E. & Shephard, N. (2004), 'Power and bipower variation with stochastic volatility and jumps', *Journal of financial econometrics* **2**(1), 1–37.

Bollerslev, T. (1986), 'Generalized autoregressive conditional heteroskedasticity', *Journal of econometrics* **31**(3), 307–327.

Bollerslev, T. (2022), 'Realized semi (co) variation: Signs that all volatilities are not created equal', *Journal of Financial Econometrics* **20**(2), 219–252.

Bollerslev, T., Hood, B., Huss, J. & Pedersen, L. H. (2018), 'Risk everywhere: Modeling and managing volatility', *The Review of Financial Studies* **31**(7), 2729–2773.

Bollerslev, T., Patton, A. J. & Quaedvlieg, R. (2016), 'Exploiting the errors: A simple approach for improved volatility forecasting', *Journal of Econometrics* **192**(1), 1–18.

Calvet, L. E., Fisher, A. J. & Thompson, S. B. (2006), 'Volatility comovement: a multifrequency approach', *Journal of econometrics* **131**(1-2), 179–215.

Chen, A. Y. & Zimmermann, T. (2022), 'Open source cross-sectional asset pricing', *Critical Finance Review* **11**(2), 207–264.

Chinco, A., Clark-Joseph, A. D. & Ye, M. (2019), 'Sparse signals in the cross-section of returns', *The Journal of Finance* **74**(1), 449–492.

Christensen, K., Oomen, R. C. & Podolskij, M. (2014), 'Fact or friction: Jumps at ultra high frequency', *Journal of Financial Economics* **114**(3), 576–599.

Corsi, F. (2009), 'A simple approximate long-memory model of realized volatility', *Journal of Financial Econometrics* **7**(2), 174–196.

Corsi, F., Pirino, D. & Renò, R. (2010), 'Threshold bipower variation and the impact of jumps on volatility forecasting', *Journal of Econometrics* **159**(2), 276–288.

Davidson, J. (1994), *Stochastic Limit Theory*, Oxford University Press.

Diebold, F. X. & Mariano, R. S. (2002), 'Comparing predictive accuracy', *Journal of Business & economic statistics* **20**(1), 134–144.

Ding, Y., Engle, R., Li, Y. & Zheng, X. (2025), 'Multiplicative factor model for volatility', *Journal of Econometrics* **249**, 105959.

Engle, R. F. (1982), 'Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation', *Econometrica: Journal of the econometric society* pp. 987–1007.

Engle, R. F. & Campos-Martins, S. (2023), 'What are the events that shake our world? measuring and hedging global covol', *Journal of Financial Economics* **147**(1), 221–242.

Engle, R. F., Ito, T. & Lin, W.-L. (1988), 'Meteor showers or heat waves? heteroskedastic intra-daily volatility in the foreign exchange market'.

Fama, E. F. & French, K. R. (2015), 'A five-factor asset pricing model', *Journal of financial economics* **116**(1), 1–22.

Fama, E. F. & French, K. R. (2018), 'Choosing factors', *Journal of financial economics* **128**(2), 234–252.

Freyberger, J., Neuhierl, A. & Weber, M. (2020), 'Dissecting characteristics nonparametrically', *The Review of Financial Studies* **33**(5), 2326–2377.

Gençay, R., Dacorogna, M., Muller, U. A., Pictet, O. & Olsen, R. (2001), *An introduction to high-frequency finance*, Elsevier.

Gu, S., Kelly, B. & Xiu, D. (2020), 'Empirical asset pricing via machine learning', *The Review of Financial Studies* **33**(5), 2223–2273.

Harvey, C. R., Liu, Y. & Zhu, H. (2016), '... and the cross-section of expected returns', *The Review of Financial Studies* **29**(1), 5–68.

Herskovic, B., Kelly, B., Lustig, H. & Van Nieuwerburgh, S. (2016), 'The common factor in idiosyncratic volatility: Quantitative asset pricing implications', *Journal of Financial Economics* **119**(2), 249–283.

Herskovic, B., Kelly, B., Lustig, H. & Van Nieuwerburgh, S. (2020), 'Firm volatility in granular networks', *Journal of Political Economy* **128**(11), 4097–4162.

Hizmeri, R., Izzeldin, M., Nolte, I. & Pappas, V. (2022), 'A generalized heterogeneous autoregressive model using market information', *Quantitative Finance* **22**(8), 1513–1534.

Holden, C. W. & Jacobsen, S. (2014), 'Liquidity measurement problems in fast, competitive markets: Expensive and cheap solutions', *The Journal of Finance* **69**(4), 1747–1785.

Hong, S. Y., Nolte, I., Taylor, S. J. & Zhao, X. (2023), 'Volatility estimation and forecasts based on price durations', *Journal of Financial Econometrics* **21**(1), 106–144.

Hou, K., Xue, C. & Zhang, L. (2018), 'Replicating anomalies', *The Review of Financial Studies* **33**(5), 2019–2133.

Jacod, J. & Shiryaev, A. N. (2003), *Limit Theorems for Stochastic Processes*, Springer.

Jensen, T. I., Kelly, B. & Pedersen, L. H. (2023), 'Is there a replication crisis in finance?', *The Journal of Finance* **78**(5), 2465–2518.

Kapadia, N., Linn, M. & Paye, B. (2024), 'One vol to rule them all: Common volatility dynamics in factor returns', *Journal of Financial and Quantitative Analysis* **59**(3), 1185–1212.

Li, Y., Nolte, I., Nolte, S. & Yu, S. (2025), 'Realized candlestick wicks', *Journal of Econometrics* **250**, 106014.

Luciani, M. & Veredas, D. (2015), 'Estimating and forecasting large panels of volatilities with approximate dynamic factor models', *Journal of Forecasting* **34**(3), 163–176.

Mancini, C. (2009), 'Non-parametric threshold estimation for models with stochastic diffusion coefficient and jumps', *Scandinavian Journal of Statistics* **36**(2), 270–296.

Nolte, I. & Voev, V. (2012), 'Least squares inference on integrated volatility and the relationship between efficient prices and noise', *Journal of Business & Economic Statistics* **30**(1), 94–108.

Patton, A. J. (2011), 'Volatility forecast comparison using imperfect volatility proxies', *Journal of Econometrics* **160**(1), 246–256.

Patton, A. J. & Sheppard, K. (2015), 'Good volatility, bad volatility: Signed jumps and the persistence of volatility', *Review of Economics and Statistics* **97**(3), 683–697.

Ross, S. A. (2013), The arbitrage theory of capital asset pricing, *in* 'Handbook of the fundamentals of financial decision making: Part I', World Scientific, pp. 11–30.

Sharpe, W. F. (1964), 'Capital asset prices: A theory of market equilibrium under conditions of risk', *The journal of finance* **19**(3), 425–442.

Taylor, S. J. (1982), 'Financial returns modelled by the product of two stochastic processes-a study of the daily sugar prices 1961-75', *Time series analysis: theory and practice* **1**, 203–226.

Zhang, L., Mykland, P. A. & Aït-Sahalia, Y. (2005), 'A tale of two time scales: Determining integrated volatility with noisy high-frequency data', *Journal of the American Statistical Association* **100**(472), 1394–1411.

Zheng, X. & Li, Y. (2011), 'On the estimation of integrated covariance matrices of high dimensional diffusion processes', *The Annals of Statistics* **39**(6), 3121–3151.