

Matrix-based Prediction Approach for Intraday Instantaneous Volatility Vector

Sung Hoon Choi* Donggyu Kim†

December 26, 2024

Abstract

In this paper, we introduce a novel method for predicting intraday instantaneous volatility based on Itô semimartingale models using high-frequency financial data. Several studies have highlighted stylized volatility time series features, such as interday auto-regressive dynamics and the intraday U-shaped pattern. To accommodate these volatility features, we propose an interday-by-intraday instantaneous volatility matrix process that can be decomposed into low-rank conditional expected instantaneous volatility and noise matrices. To predict the low-rank conditional expected instantaneous volatility matrix, we propose the Two-Side Projected-PCA (TIP-PCA) procedure. We establish asymptotic properties of the proposed estimators and conduct a simulation study to assess the finite sample performance of the proposed prediction method. Finally, we apply the TIP-PCA method to an out-of-sample instantaneous volatility vector prediction study using high-frequency data from the S&P 500 index and 11 sector index funds.

Key words: Diffusion process, high-frequency financial data, low-rank matrix, semi-parametric factor models.

*Department of Economics, University of Connecticut, Storrs, CT 06269, USA. E-mail: sung_hoon.choi@uconn.edu.

†Department of Economics, University of California, Riverside, CA 92521, USA. Email: donggyu.kim@ucr.edu.

1 Introduction

The analysis of volatility is a vibrant research area in financial econometrics and statistics. In practice, it is crucial to investigate the volatility dynamics of asset returns for hedging, option pricing, risk management, and portfolio management. With the wide availability of high-frequency financial data, several well-performing non-parametric integrated volatility estimation methods have been developed. Examples include two-time scale realized volatility (TSRV) (Zhang et al., 2005), multi-scale realized volatility (MSRV) (Zhang, 2006, 2011), pre-averaging realized volatility (PRV) (Christensen et al., 2010; Jacod et al., 2009), wavelet realized volatility (WRV) (Fan and Wang, 2007), kernel realized volatility (KRV) (Barndorff-Nielsen et al., 2008, 2011), quasi-maximum likelihood estimator (QMLE) (Aït-Sahalia et al., 2010; Xiu, 2010), local method of moments (Bibinger et al., 2014), and robust pre-averaging realized volatility (Fan and Kim, 2018; Shin et al., 2023). This incorporation of high-frequency information enhances our understanding of low-frequency (i.e., interday) market dynamics, and several conditional volatility models have been developed based on the realized volatility to explain market dynamics. Examples include realized volatility-based modeling approaches (Andersen et al., 2003), heterogeneous auto-regressive (HAR) models (Corsi, 2009), high-frequency-based volatility (HEAVY) models (Shephard and Sheppard, 2010), realized GARCH models (Hansen et al., 2012), and unified GARCH-Itô models (Kim and Wang, 2016; Song et al., 2021).

To understand intraday dynamics, several non-parametric instantaneous (or spot) volatility estimation procedures have been developed (Fan and Wang, 2008; Figueroa-López and Wu, 2022; Foster and Nelson, 1996; Kristensen, 2010; Mancini et al., 2015; Todorov, 2019; Todorov and Zhang, 2023; Zu and Boswijk, 2014). With these well-performing instantaneous volatility estimators, several studies have found that intraday instantaneous volatility has U-shaped patterns (Admati and Pfleiderer, 1988; Andersen and Bollerslev, 1997; Andersen et al., 2019; Hong and Wang, 2000; Li and Linton, 2023). On the other hand, forecasting intraday volatility has not received as much attention in the literature compared to daily

volatility forecasting. [Engle and Sokalska \(2012\)](#) developed a modified GARCH model that imposes the conditional variance as a product of daily, diurnal, and stochastic intraday components for intraday volatility forecasting. Recently, [Zhang et al. \(2024\)](#) studied several non-parametric machine learning methods to forecast intraday realized volatility by utilizing commonality in intraday volatility. From the aforementioned studies, it is evident that interday volatility dynamics can be explained by autoregressive-type time series dynamics while intraday volatility dynamics have some periodic patterns, such as a U-shape. Thus, to predict the one-day-ahead whole intraday instantaneous volatility vector, we need to consider the interday and intraday dynamics simultaneously. Furthermore, since we consider the prediction of whole intraday instantaneous volatilities, we need to handle the overparameterization issue.

This paper introduces a novel approach for predicting the one-day-ahead instantaneous volatility process. Specifically, we represent the instantaneous volatility process in a matrix form. For example, to account for interday time series dynamics and intraday periodic patterns, each row corresponds to a day, and each column represents a high-frequency sequence within that day. Thus, we have an interday-by-intraday instantaneous volatility matrix. To handle the overparameterization issue, we impose a low-rank plus noise structure on the instantaneous volatility matrix, where the low-rank components represent a conditional expected instantaneous volatility matrix and have semiparametric factor structures. To accommodate the proposed instantaneous volatility model, we adopt the Projected-PCA method ([Fan et al., 2016b](#)) to estimate left-singular and right-singular vector components with instantaneous volatility matrix estimators. For example, for the left-singular vector, we project the left-singular vectors onto a linear space spanned by past realized volatility estimators, which enables us to account for the interday time-series dynamics and to predict the one-day-ahead instantaneous volatility vector with observed current realized volatility estimators. Conversely, for the right-singular vector, to explain periodic patterns, we project the right-singular vectors onto a linear space spanned by deterministic time sequences. This

two-side projection enables us to explain the interday and intraday dynamics simultaneously. We call this the Two-Side Projected-PCA (TIP-PCA) procedure. We then derive convergence rates for the projected instantaneous volatility matrix estimator and the predicted instantaneous volatility vector using the TIP-PCA method. As an empirical study, we apply the proposed TIP-PCA estimator on out-of-sample predictions for the one-day-ahead instantaneous volatility process using high-frequency trading data. Using a rolling window scheme, we compare TIP-PCA with several methods. TIP-PCA demonstrates superior performance by leveraging both interday and intraday dynamics. Additionally, we evaluate one-day-ahead 10-minute frequency Value at Risk (VaR) predictions, where TIP-PCA consistently outperforms other methods in backtests.

The remainder of the paper is structured as follows. Section 2 establishes the model and introduces the TIP-PCA prediction procedure. Section 3 provides an asymptotic analysis of the TIP-PCA estimators. The effectiveness of the proposed method is demonstrated through a simulation study in Section 4 and by applying it to real high-frequency financial data for predicting the one-day-ahead instantaneous volatility process in Section 5. Section 6 concludes the study. All proofs are presented in the online supplementary file.

2 Model Setup and Estimation Procedure

Throughout this paper, we denote by $\|\mathbf{A}\|_F$, $\|\mathbf{A}\|_2$ (or $\|\mathbf{A}\|$ for short), $\|\mathbf{A}\|_1$, $\|\mathbf{A}\|_\infty$, and $\|\mathbf{A}\|_{\max}$ the Frobenius norm, operator norm, l_1 -norm, l_∞ -norm, and elementwise norm, which are defined, respectively, as $\|\mathbf{A}\|_F = \text{tr}^{1/2}(\mathbf{A}'\mathbf{A})$, $\|\mathbf{A}\|_2 = \lambda_{\max}^{1/2}(\mathbf{A}'\mathbf{A})$, $\|\mathbf{A}\|_1 = \max_j \sum_i |a_{ij}|$, $\|\mathbf{A}\|_\infty = \max_i \sum_j |a_{ij}|$, and $\|\mathbf{A}\|_{\max} = \max_{i,j} |a_{ij}|$. When \mathbf{A} is a vector, the maximum norm is denoted as $\|\mathbf{A}\|_\infty = \max_i |a_i|$, and both $\|\mathbf{A}\|$ and $\|\mathbf{A}\|_F$ are equal to the Euclidean norm.

2.1 A Model Setup

We consider the following jump diffusion process: for the i -th day and intraday time $t \in [0, 1]$,

$$dX_{i,t} = \mu_{i,t}dt + \sigma_{i,t}dB_{i,t} + J_{i,t}dP_{i,t}, \quad (2.1)$$

where $X_{i,t}$ is the log price of an asset, $\mu_{i,t}$ is a drift process, $B_{i,t}$ is a one-dimensional standard Brownian motion, $J_{i,t}$ is the jump size, and $P_{i,t}$ is the Poisson process with the intensity $\mu_{i,t}$. For a given intraday time sequence, for each $i = 1, \dots, D$ and $j = 1, \dots, n$, we denote the instantaneous volatility process as $c_{i,j} := \sigma_{i,t_j}^2$, where $0 < t_1 < \dots < t_n = 1$. Then, we can write the discrete-time instantaneous volatility process as follows:

$$\Sigma_{D,n} = (c_{i,j})_{D \times n} = \mathbf{U}\mathbf{\Lambda}\mathbf{V}' + \Sigma_\varepsilon, \quad (2.2)$$

where $\mathbf{U} = (u_{i,k})_{i=1,\dots,D,k=1,\dots,r}$ is the left singular vector matrix, $\mathbf{V} = (v_{j,k})_{j=1,\dots,n,k=1,\dots,r}$ is the right singular vector matrix, $\mathbf{\Lambda} = \text{Diag}(\lambda_1, \dots, \lambda_r)$ is the singular value matrix, and $\Sigma_\varepsilon = (\varepsilon_{i,j})_{D \times n}$ is the random noise matrix. We note that $\Sigma_{D,n}$ is a $D \times n$ matrix, which is distinct from a covariance matrix and contains only positive elements. The left singular vector matrix represents interday volatility dynamics, while the right singular vector matrix explains intraday volatility dynamics. For example, we consider $r = 1$. For the intraday, the U-shaped instantaneous volatility pattern is often observed in empirical data and supported by the financial market ([Admati and Pfleiderer, 1988](#); [Andersen and Bollerslev, 1997](#); [Andersen et al., 2019](#); [Hong and Wang, 2000](#)). Thus, we can use a U-shape function with respect to time t for \mathbf{V} , for example, $v_{t,1} = a_1(t - a_2)^2 + a_3$. In contrast, for the interday, the daily dynamics are often explained by past realized volatilities ([Corsi, 2009](#); [Hansen et al., 2012](#); [Kim and Fan, 2019](#); [Kim and Wang, 2016](#); [Shephard and Sheppard, 2010](#); [Song et al., 2021](#)). To reflect this, $u_{i,1}$ is a function of past realized volatilities, such as the HAR model ([Corsi, 2009](#)), for example, $u_{i,1} = b_0 + b_1RV_{i-1} + b_2\frac{1}{5}\sum_{j=1}^5RV_{i-j} + b_3\frac{1}{22}\sum_{j=1}^{22}RV_{i-j}$, where RV_i is the i -th day

realized volatility. The features mentioned above motivate the representation of the model (2.2). We note that in this paper, to simplify prediction modeling, we consider the time series structure of interday volatility dynamics and the instantaneous periodic pattern of intraday volatility dynamics, as described in the example above. Our empirical experiment supports that the simplified structure performs well. However, we can impose more complicated structures on both dimensions. That is, the above example is one of the possible models, and we propose a generalized model discussed below.

In this paper, our goal is to predict the instantaneous volatility process for the next day. In general, we assume that $\lambda_1, \dots, \lambda_r$ are latent factors, $u_{i,k}$ is \mathcal{F}_{i-1} -adapted, and $v_{t,k}$ is a function of time t . Thus, given \mathcal{F}_D , we can predict the instantaneous volatility as follows:

$$E[\sigma_{D+1,t}^2 | \mathcal{F}_D] = \sum_{k=1}^r \lambda_k u_{D+1,k} v_{t,k} \text{ a.s.} \quad (2.3)$$

Given (2.2), we impose the following nonparametric structure on both singular vectors: for each $k \leq r$, $i \leq D$, and $j \leq n$,

$$u_{i,k} = g_k(\mathbf{x}_i), \quad v_{j,k} = h_k(\mathbf{w}_j), \quad (2.4)$$

where $\mathbf{x}_i = (x_{i1}, \dots, x_{id_1})$ and $\mathbf{w}_j = (w_{j1}, \dots, w_{jd_2})$ are observable covariates that explain the left and right singular vectors, respectively. In this context, \mathbf{x}_i can be the past realized volatility of yesterday, last week, and last month in the HAR model, while \mathbf{w}_j can be the intraday time sequence. Furthermore, we assume that each unknown nonparametric function is additive as follows: for each $k \leq r$, $i \leq D$, and $j \leq n$,

$$\begin{aligned} g_k(\mathbf{x}_i) &= \phi(\mathbf{x}_i)' \mathbf{b}_k + R_k(\mathbf{x}_i), \\ h_k(\mathbf{w}_j) &= \psi(\mathbf{w}_j)' \mathbf{a}_k + Q_k(\mathbf{w}_j), \end{aligned}$$

where $\phi(\mathbf{x}_i)$ is a $(J_1 d_1) \times 1$ vector of basis functions, \mathbf{b}_k is a $(J_1 d_1) \times 1$ vector of sieve

coefficients, and $R_k(\mathbf{x}_i)$ is the approximation error term; $\psi(\mathbf{w}_j)$ is a $(J_2 d_2) \times 1$ vector of basis functions, \mathbf{a}_k is a $(J_2 d_2) \times 1$ vector of sieve coefficients, and $Q_k(\mathbf{w}_j)$ is the approximation error term. Equivalently, we can write $g_k(\mathbf{x}_i) = \sum_{d=1}^{d_1} g_{kd}(x_{id})$ and $h_k(\mathbf{w}_j) = \sum_{d=1}^{d_2} h_{kd}(w_{jd})$, where $g_{kd}(x_{id}) = \sum_{l=1}^{J_1} b_{l,kd} \phi_l(x_{id}) + R_{kd}(x_{id})$ and $h_{kd}(w_{jd}) = \sum_{l=1}^{J_2} a_{l,kd} \psi_l(w_{jd}) + Q_{kd}(w_{jd})$, respectively. Hence, each additive component of g_k and h_k can be estimated by the sieve method. Throughout the paper, we assume that $d_1 = \dim(\mathbf{x}_i)$, $d_2 = \dim(\mathbf{w}_j)$ and r are fixed. The number of sieve terms, J_1 and J_2 , grow very slowly as $D \rightarrow \infty$ and $n \rightarrow \infty$, respectively. In a matrix form, we can write

$$\begin{aligned}\mathbf{U} &:= \mathbf{G}(\mathbf{X}) = \Phi(\mathbf{X})\mathbf{B} + \mathbf{R}(\mathbf{X}), \\ \mathbf{V} &:= \mathbf{H}(\mathbf{W}) = \Psi(\mathbf{W})\mathbf{A} + \mathbf{Q}(\mathbf{W}),\end{aligned}$$

where the $D \times (J_1 d_1)$ matrix $\Phi(\mathbf{X}) = (\phi(\mathbf{x}_1), \dots, \phi(\mathbf{x}_D))'$, the $(J_1 d_1) \times r$ matrix $\mathbf{B} = (\mathbf{b}_1, \dots, \mathbf{b}_r)$, and $\mathbf{R}(\mathbf{X}) = (R_k(\mathbf{x}_i))_{D \times r}$; the $n \times (J_2 d_2)$ matrix $\Psi(\mathbf{W}) = (\psi(\mathbf{w}_1), \dots, \psi(\mathbf{w}_n))'$, the $(J_2 d_2) \times r$ matrix $\mathbf{A} = (\mathbf{a}_1, \dots, \mathbf{a}_r)$, and $\mathbf{Q}(\mathbf{W}) = (Q_k(\mathbf{w}_j))_{n \times r}$. We note that since $\mathbf{G}(\mathbf{X})$ and $\mathbf{H}(\mathbf{W})$ are identifiable up to the sign and $\Phi(\mathbf{X})$ and $\Psi(\mathbf{W})$ are given, the coefficients \mathbf{A} and \mathbf{B} can be defined by the definition of linear regression coefficients. Then, since the vectors of basis functions are non-singular, we can define uniquely $\Phi(\mathbf{X})\mathbf{B}$ and $\Psi(\mathbf{W})\mathbf{A}$.

Due to the imperfections of the trading mechanisms (Aït-Sahalia and Yu, 2009), the true underlined log-stock price $X_{i,t}$ in (2.1) is not observable. To reflect the imperfections, we assume that the high-frequency intraday observations X_{i,t_s} , $s = 1, \dots, m$, are contaminated by microstructure noises as follows:

$$Y_{i,t_s} = X_{i,t_s} + e_{i,t_s}, \quad i = 1, \dots, D, s = 1, \dots, m, \quad (2.5)$$

where the microstructure noises e_{i,t_s} are independent random variables with a mean of zero and a variance of η_{ii} . For simplicity, we assume that the observed time points are equally spaced, that is, $t_s - t_{s-1} = m^{-1}$ for $i = 1, \dots, D$ and $s = 2, \dots, m$.

Several non-parametric instantaneous volatility estimation procedures have been developed (Fan and Wang, 2008; Figueroa-López and Wu, 2022; Foster and Nelson, 1996; Kristensen, 2010; Mancini et al., 2015; Todorov, 2019; Todorov and Zhang, 2023; Zu and Boswijk, 2014). We can use any well-performing instantaneous volatility estimator that satisfies Assumption 3.1 (ii). In the numerical study, we employ the jump robust pre-averaging method proposed by Figueroa-López and Wu (2022). The specific method is described in (4.1).

2.2 Two-Side Projected-PCA

To accommodate the semiparametric structure in Section 2.1, we need to project the left and right singular vectors on linear spaces spanned by the corresponding covariates. To do this, we apply the Projected-PCA (Fan et al., 2016b) procedure to the left and right singular vectors with the well-performing instantaneous volatility estimator. The specific procedure is as follows:

1. For each $i \leq D$ and $j \leq n$, we estimate the instantaneous volatility, $c_{i,j} = \sigma_{i,t_j}^2$, using high-frequency log-price observations and denote them $\hat{c}_{i,j}$. Let $\hat{\mathbf{\Lambda}} = \text{diag}(\hat{\lambda}_1, \dots, \hat{\lambda}_r)$, where $\hat{\lambda}_1 \geq \hat{\lambda}_2 \geq \dots \geq \hat{\lambda}_r$ are the square root of the leading eigenvalues of $\hat{\Sigma}_{D,n} \hat{\Sigma}'_{D,n}$, where $\hat{\Sigma}_{D,n} = (\hat{c}_{i,j})_{D \times n}$.
2. Define the $D \times D$ projection matrix as $\mathbf{P}_\Phi = \Phi(\mathbf{X})(\Phi(\mathbf{X})'\Phi(\mathbf{X}))^{-1}\Phi(\mathbf{X})'$. The columns of $\hat{\mathbf{G}}(\mathbf{X}) := (\hat{U}_1, \dots, \hat{U}_r)$ are defined as the r leading eigenvectors of the $D \times D$ matrix $\mathbf{P}_\Phi \hat{\Sigma}_{D,n} \hat{\Sigma}'_{D,n} \mathbf{P}_\Phi$. Then, we can estimate \mathbf{B} by

$$\hat{\mathbf{B}} = (\hat{\mathbf{b}}_1, \dots, \hat{\mathbf{b}}_r) = (\Phi(\mathbf{X})'\Phi(\mathbf{X}))^{-1}\Phi(\mathbf{X})'\hat{\mathbf{G}}(\mathbf{X}).$$

Given any $\mathbf{x} \in \mathcal{X}$, we estimate $g_k(\cdot)$ by

$$\hat{g}_k(\mathbf{x}) = \phi(\mathbf{x})'\hat{\mathbf{b}}_k \quad \text{for } k = 1, \dots, r,$$

where \mathcal{X} denotes the support of \mathbf{x}_i .

3. Define the $n \times n$ projection matrix as $\mathbf{P}_\Psi = \Psi(\mathbf{W})(\Psi(\mathbf{W})'\Psi(\mathbf{W}))^{-1}\Psi(\mathbf{W})'$. The columns of $\hat{\mathbf{H}}(\mathbf{W}) := (\hat{V}_1, \dots, \hat{V}_r)$ are defined as the r leading eigenvectors of the $n \times n$ matrix $\mathbf{P}_\Psi \hat{\Sigma}'_{D,n} \hat{\Sigma}_{D,n} \mathbf{P}_\Psi$.
4. We estimate a sign vector $s_0 = (s_{01}, \dots, s_{0r}) \in \{-1, 1\}^r$ defined in (2.7) by

$$\hat{s} := (\hat{s}_1, \dots, \hat{s}_r) = \underset{s \in \{-1, 1\}^r}{\operatorname{argmin}} \left\| \sum_{k=1}^r s_k \hat{\lambda}_k \hat{U}_k \hat{V}_k' - \hat{\Sigma}_{D,n} \right\|_F^2. \quad (2.6)$$

Then, we update the right singular vector estimator by $\hat{\mathbf{H}}(\mathbf{W}) = (\hat{s}_1 \hat{V}_1, \dots, \hat{s}_r \hat{V}_r)$.

5. Finally, we predict the conditional expectation of the one-day-ahead instantaneous volatility vector $E[(c_{D+1,1}, \dots, c_{D+1,n}) | \mathcal{F}_D]$ by

$$(\tilde{c}_{D+1,1}, \dots, \tilde{c}_{D+1,n}) = \hat{\mathbf{g}}(\mathbf{x}_{D+1}) \hat{\Lambda} \hat{\mathbf{H}}(\mathbf{W})',$$

where $\hat{\mathbf{g}}(\mathbf{x}_{D+1}) = (\hat{g}_1(\mathbf{x}_{D+1}), \dots, \hat{g}_r(\mathbf{x}_{D+1}))$.

Remark 2.1. To estimate the instantaneous volatility matrix, we need to match the signs for singular vector estimators (Cho et al., 2017). This is because \hat{U}_k and \hat{V}_k can estimate $U_k = (u_{1,k}, \dots, u_{D,k})'$ and $V_k = (v_{1,k}, \dots, v_{D,k})'$ up to signs, such as $\operatorname{sign}(\langle \hat{U}_k, U_k \rangle)$ and $\operatorname{sign}(\langle \hat{V}_k, V_k \rangle)$, respectively. Let $s_0 = (s_{01}, \dots, s_{0r}) \in \{-1, 1\}^r$ be

$$s_{0k} = \operatorname{sign}(\langle \hat{U}_k, U_k \rangle) \operatorname{sign}(\langle \hat{V}_k, V_k \rangle) \quad \text{for } k = 1, \dots, r. \quad (2.7)$$

Then, $\mathbf{U}\Lambda\mathbf{V}'$ can be consistently estimated by $\sum_{k=1}^r s_{0k} \hat{\lambda}_k \hat{U}_k \hat{V}_k'$. However, since s_0 is unknown in practice, we employ the sign estimation procedure as discussed in Step 2.6 above. We can show their sign consistency under a regularity condition (see Assumption 3.5).

In summary, given the estimated instantaneous volatility matrix, we initially estimate

the singular values using the conventional PCA method. We employ the Projected-PCA method (Fan et al., 2016b) to estimate the unknown nonparametric function using observable covariates (e.g., a series of past realized volatilities). We then apply the Projected-PCA method again to estimate the right singular vector matrix using the covariate (e.g., intraday time sequence). Finally, with the observable covariates \mathbf{x}_{D+1} (i.e., information about past realized volatilities on the D th day), we predict the one-day-ahead instantaneous volatility process by multiplying the estimated singular value and vector components. We refer to this procedure as the Two-side Projected-PCA (TIP-PCA). The TIP-PCA method can accurately predict instantaneous volatility by incorporating both interday and intraday dynamics, as the projection approach removes noise components. Moreover, since we assume a low-rank matrix, we can reduce the complexity of the model, which helps overcome the overparameterization. The numerical study in Sections 4 and 5 demonstrates that TIP-PCA performs well in predicting a one-day-ahead instantaneous volatility vector.

2.3 Choice of Tuning Parameters

The suggested TIP-PCA estimator requires the choice of tuning parameters r , J_1 , and J_2 . First, the number of latent factors, r , can be chosen through data-driven methods (Ahn and Horenstein, 2013; Bai and Ng, 2002; Onatski, 2010). For example, r can be determined by finding the largest singular value gap or singular value ratio such that $\max_{k \leq r_{\max}} (\hat{\lambda}_k - \hat{\lambda}_{k+1})$ and $\max_{k \leq r_{\max}} \frac{\hat{\lambda}_k}{\hat{\lambda}_{k+1}}$ for a predetermined maximum number of factors r_{\max} . For the numerical studies in Sections 4 and 5, we employed rank 1 using the eigenvalue ratio method proposed by Ahn and Horenstein (2013).

The numbers of sieve terms, J_1 and J_2 , and basis functions can be flexibly chosen by practitioners based on the conjecture of the nonparametric function form (Chen et al., 2020; Fan et al., 2016b). We note that J_1 and J_2 are the cost to approximate the unknown nonparametric functions $g_k(\cdot)$ and $h_k(\cdot)$ (see Remark 3.2). Both J_1 and J_2 grow very slowly such as $\log D$ and $\log n$, respectively, and we require $J_1 = o(\sqrt{D})$ and $J_2 = o(\sqrt{n})$ as described

in Proposition 3.1 and Theorem 3.1. In this context, the interday volatility dynamic is a linear function of the past realized volatilities, while the intraday volatility dynamic is a U-shaped function with respect to deterministic time sequences. Therefore, for the numerical studies in Sections 4 and 5, we employed the additive polynomial basis with the sieve dimensions $J_1 = 2$ and $J_2 = 3$ for the TIP-PCA method.

3 Asymptotic Properties

In this section, we establish the asymptotic properties of the proposed TIP-PCA estimator. To do this, we impose the following technical assumptions.

Assumption 3.1.

(i) For $k \leq r$, the eigengap satisfies $|\lambda_{k+1} - \lambda_k| = O_P(\sqrt{nD})$ and $\lambda_{r+1} = 0$.

(ii) For each $i \leq D$ and $j \leq n$ and, the estimated instantaneous volatility $\widehat{c}_{i,j}$ satisfies

$$\widehat{c}_{i,j} - c_{i,j} = v_{i,j} + \varsigma_{i,j},$$

where $v_{i,j}$ follows the martingale difference sequence and $\varsigma_{i,j}$ is the estimation bias term such that $E(v_{i,j}|\mathcal{F}_{i,t_j}) = 0$ a.s., $E(v_{i,j}^6) = O(m^{-3/4})$, $E(\varsigma_{i,j}^6) = O(\rho_m^6)$, and for any $s \leq D$ and $j' \leq n$, $E(c_{i,j'}v_{i,j}|\mathcal{F}_{i,t_j}) = 0$ and $E(c_{s,j}c_{s,j'}v_{i,j}|\mathcal{F}_{i,t_j}) = 0$ a.s. Additionally, there exist a constant $C_1 > 0$ such that $E(c_{i,j}^4) < C_1$.

(iii) Let $\mathbf{\Omega}_{\varepsilon, D \times D}$ denote the $D \times D$ covariance matrix of $\varepsilon_j = (\varepsilon_{1,j}, \dots, \varepsilon_{D,j})'$. Define $\mathbf{\Omega}_{D \times D} = \frac{1}{n} \mathbf{U} \mathbf{\Lambda}^2 \mathbf{U}' + \mathbf{\Omega}_{\varepsilon, D \times D}$. The instantaneous volatility matrix, $\mathbf{\Sigma}_{D,n}$, satisfies

$$\left\| \frac{1}{n} \mathbf{\Sigma}_{D,n} \mathbf{\Sigma}'_{D,n} - \mathbf{\Omega}_{D \times D} \right\|_{\max} = O_P(\sqrt{\log D/n}).$$

Assumption 3.1 is related to assumptions for the instantaneous volatility matrix. Assumption 3.1(i) is the eigengap assumption, which is essential for analyzing low-rank ma-

trices (Candes and Plan, 2010; Cho et al., 2017; Fan et al., 2018a). We note that since we have a $D \times n$ instantaneous volatility matrix, the pervasive condition implies that the eigenvalue for the low-rank component has \sqrt{nD} order. The moment conditions in Assumption 3.1(ii) can be satisfied under some assumptions on the process X , microstructure noise, and kernel function in Figueroa-López and Wu (2022) by using similar arguments in Kim et al. (2016). When estimating the spot volatility at time t , we usually use data after time t . This implies $E(v_{i,j}|\mathcal{F}_{i,t_j}) = 0$ a.s. We note that the spot volatility estimator is asymptotically unbiased, thus, the estimation bias term goes to zero faster than $m^{-\frac{1}{8}}$. The uncorrelated condition, $E(c_{i,j}'v_{i,j}|\mathcal{F}_{i,t_j}) = E(c_{s,j}c_{s,j}'v_{i,j}|\mathcal{F}_{i,t_j}) = 0$ a.s., is required to show the benefit of the smoothing scheme. This condition is satisfied, if the spot volatility can be represented by \mathcal{F}_{i,t_j} -adapted processes and independent noises. For example, the usual time series structure, such as ARMA, satisfies it. Thus, this condition is not restrictive. It is worth noting that if the uncorrelated condition is not satisfied, we have a slower convergence rate, for example, we have additionally $m^{-1/8}$ in Proposition 3.1 and Theorem 3.1. To analyze large matrix inferences, we impose the element-wise convergence condition (Assumption 3.1(iii)). This condition can be easily satisfied under the sub-Gaussian condition and the mixing time dependency (Fan et al., 2018a,b; Vershynin, 2010; Wang and Fan, 2017). We can also obtain this condition under heavy-tailed observations with the bounded fourth moments (Fan and Kim, 2018; Shin et al., 2023).

Assumption 3.2.

- (i) There are c_{\min} and $c_{\max} > 0$ so that, with the probability approaching one, as $D \rightarrow \infty$ and $n \rightarrow \infty$,

$$c_{\min} < \lambda_{\min}(D^{-1}\Phi(\mathbf{X})'\Phi(\mathbf{X})) < \lambda_{\max}(D^{-1}\Phi(\mathbf{X})'\Phi(\mathbf{X})) < c_{\max},$$

$$c_{\min} < \lambda_{\min}(n^{-1}\Psi(\mathbf{W})'\Psi(\mathbf{W})) < \lambda_{\max}(n^{-1}\Psi(\mathbf{W})'\Psi(\mathbf{W})) < c_{\max}.$$

- (ii) $\max_{l \leq J_1, i \leq D, d \leq d_1} E\phi_l(x_{id})^2 < \infty$, and $\max_{l \leq J_2, j \leq n, d \leq d_2} E\psi_l(w_{jd})^2 < \infty$.

(iii) $\max_{k \leq r, i \leq D} E g_k(\mathbf{x}_i)^2 < \infty$, and $\max_{k \leq r, j \leq n} E h_k(\mathbf{w}_j)^2 < \infty$.

Assumption 3.2 is related to basis functions. Intuitively, the strong law of large numbers implies Assumption 3.2 (i), which can be satisfied by normalizing common basis functions such as B-splines, polynomial series, or Fourier basis.

Assumption 3.3. For all $d \leq d_1, d' \leq d_2, k \leq r$,

(i) the functions $g_{kd}(\cdot)$ and $h_{kd'}(\cdot)$ belong to a Hölder class \mathcal{G} and \mathcal{H} defined by, for some $L > 0$,

$$\begin{aligned}\mathcal{G} &= \{g : |g^{(c)}(s) - g^{(c)}(t)| \leq L|s - t|^\alpha\}, \\ \mathcal{H} &= \{h : |h^{(c')}(s) - h^{(c')}(t)| \leq L|s - t|^{\alpha'}\}.\end{aligned}$$

(ii) the sieve coefficients $\{b_{l,kd}\}_{l \leq J_1}$ satisfy for $\kappa = 2(c + \alpha) \geq 4$, as $J_1 \rightarrow \infty$,

$$\sup_{x \in \mathcal{X}_d} \left| g_{kd}(x) - \sum_{l=1}^{J_1} b_{l,kd} \phi_l(x) \right|^2 = O(J_1^{-\kappa}/D),$$

where \mathcal{X}_d is the support of the d th element of \mathbf{x}_i . Similarly, the sieve coefficients $\{a_{l,kd'}\}_{l \leq J_2}$ satisfy for $\kappa = 2(c' + \alpha') \geq 4$, as $J_2 \rightarrow \infty$,

$$\sup_{w \in \mathcal{W}_{d'}} \left| h_{kd'}(w) - \sum_{l=1}^{J_2} a_{l,kd'} \psi_l(w) \right|^2 = O(J_2^{-\kappa}/n),$$

where $\mathcal{W}_{d'}$ is the support of the d' th element of \mathbf{w}_j .

(iii) $\max_{k \leq r, l \leq J_1, d \leq d_1} b_{l,kd}^2 = O(1/D)$ and $\max_{k \leq r, l \leq J_2, d' \leq d_2} a_{l,kd'}^2 = O(1/n)$.

Assumption 3.3 pertains to the accuracy of the sieve approximation and can be satisfied using a common basis such as a polynomial basis or B-splines (see [Chen, 2007](#)). We have the following conditions that the idiosyncratic errors are weakly dependent on both dimensions, which are commonly imposed for high-dimensional factor analysis.

Assumption 3.4.

(i) $E\varepsilon_{i,j} = 0$ for all $i \leq D, j \leq n$; $\{\varepsilon_{i,j}\}_{i \leq D, j \leq n}$ is independent of $\{\mathbf{x}_i, \mathbf{w}_j\}_{i \leq D, j \leq n}$.

(ii) There is $C > 0$ such that

$$\begin{aligned} \max_{m \leq D} \sum_{i=1}^D |E\varepsilon_{i,j}\varepsilon_{m,j}| &< C, \\ \frac{1}{nD} \sum_{i=1}^D \sum_{m=1}^D \sum_{j=1}^n \sum_{s=1}^n |E\varepsilon_{i,j}\varepsilon_{m,s}| &< C. \end{aligned}$$

(iii) The covariance matrix $\boldsymbol{\Omega}_{\varepsilon, D \times D} = \text{cov}(\boldsymbol{\varepsilon}_j) := (\omega_{\varepsilon, ij})_{D \times D}$, where $\boldsymbol{\varepsilon}_j = (\varepsilon_{1,j}, \dots, \varepsilon_{D,j})'$, satisfies

$$\|\boldsymbol{\Omega}_{\varepsilon, D \times D}\| = O(\varphi_D), \text{ where } \varphi_D = o(D).$$

Assumption 3.4(iii) is the conventional condition on the idiosyncratic covariance matrix. For example, this includes the sparsity condition, which has been considered in many applications (Boivin and Ng, 2006; Fan et al., 2016a).

Assumption 3.5. The estimated instantaneous volatility matrix $\widehat{\boldsymbol{\Sigma}}_{D,n}$ and initial estimators $\{\widehat{\lambda}_k, \widehat{U}_k, \widehat{V}_k\}_{k=1}^r$ satisfy

$$\lim_{D \rightarrow \infty, n \rightarrow \infty} \mathbb{P} \left(\min_{s \in \{-1, 1\}^r} \left\| \sum_{k=1}^r s_k \widehat{\lambda}_k \widehat{U}_k \widehat{V}_k' - \widehat{\boldsymbol{\Sigma}}_{D,n} \right\|_F^2 < \left\| \sum_{k=1}^r s_{0k} \widehat{\lambda}_k \widehat{U}_k \widehat{V}_k' - \widehat{\boldsymbol{\Sigma}}_{D,n} \right\|_F^2 \right) = 0.$$

Remark 3.1. Assumption 3.5 is related to the sign estimation in (2.6). In this paper, we conduct the singular value decomposition to estimate the left and right singular vectors separately, but these vectors can only be estimated up to a sign due to the sign problem of singular vectors. Hence, to define the signs uniquely, we impose this identifiability condition. To understand Assumption 3.5, for simplicity, we consider the case of $r = 1$. When $s_0 = 1$,

Assumption 3.5 implies

$$\lim_{D \rightarrow \infty, n \rightarrow \infty} \mathbb{P} \left(\left\| -\widehat{\lambda}_1 \widehat{U}_1 \widehat{V}_1' - \widehat{\Sigma}_{D,n} \right\|_F^2 < \left\| \widehat{\lambda}_1 \widehat{U}_1 \widehat{V}_1' - \widehat{\Sigma}_{D,n} \right\|_F^2 \right) = 0.$$

That is, the probability that \widehat{s} chooses a different sign than the true sign goes to zero as dimensions increase (Cho et al., 2017). Thus, Assumption 3.5 guarantees the identifiability of the sign problem. In light of this, Assumption 3.5 is the natural assumption to make.

We obtain the following elementwise convergence rate of the projected instantaneous volatility matrix estimator.

Proposition 3.1. *Suppose that Assumptions 3.1–3.5 hold, $J_1 = o(\sqrt{D})$, and $J_2 = o(\sqrt{n})$. As $D, n, m, J_1, J_2 \rightarrow \infty$, we have*

$$\begin{aligned} & \left\| \widehat{\mathbf{G}}(\mathbf{X}) \widehat{\Lambda} \widehat{\mathbf{H}}(\mathbf{W})' - \mathbf{G}(\mathbf{X}) \Lambda \mathbf{H}(\mathbf{W})' \right\|_{\max} \\ &= O_P \left(\sqrt{\frac{\log D}{n}} + \frac{\varphi_D}{D} + \max(J_1, J_2) \left(\sqrt{\frac{\rho_m}{m^{1/8}}} + \rho_m \right) \right. \\ & \quad \left. + \frac{J_1}{\sqrt{nm^{1/4}}} + \frac{J_2}{\sqrt{Dm^{1/4}}} + \min(J_1, J_2)^{\frac{1}{2} - \frac{\kappa}{2}} + \frac{J_1}{\sqrt{D}} + \frac{J_2}{\sqrt{n}} \right). \end{aligned}$$

Remark 3.2. Proposition 3.1 shows that the projected instantaneous volatility matrix estimator has the convergence rate $J \sqrt{\frac{\rho_m}{m^{1/8}}} + J^{\frac{1}{2} - \frac{\kappa}{2}} + \frac{J}{\sqrt{D}} + \frac{J}{\sqrt{n}}$ up to the log order and the sparsity level, when $\rho_m < m^{-1/8}$ and $J = J_1 = J_2$. We note that the convergence rate ρ_m of the bias term is faster than the convergence rate $m^{-1/8}$ of the martingale difference term in the spot volatility estimator, and the term $\sqrt{\frac{\rho_m}{m^{1/8}}}$ arises from the cross product between two components. The bias comes from the drift term that has the order of m^{-1} , and due to the subsampling scheme to handle the microstructure noise, ρ_m usually is the order of $m^{-1/2}$. The martingale term is negligible because of the smoothing effect when calculating the singular components. It is worth noting that if ρ_m is zero, with the fixed m , we can obtain the consistency. The term $D^{-1/2}$ is the cost to learn daily times series dynamics, while

the term $n^{-1/2}$ is the learning cost of the intraday periodic patterns. We note that J_1 and J_2 are related to the sieve approximation. That is, $\max(J_1, J_2)$ is the cost to approximate the unknown nonparametric functions $g_k(\cdot)$ and $h_k(\cdot)$. In addition, the term $\min(J_1, J_2)^{\frac{1}{2}-\frac{\kappa}{2}}$ is due to the approximation error. If $g_k(\cdot)$ and $h_k(\cdot)$ are known, the convergence rate is $\sqrt{\frac{\rho_m}{m^{1/8}}} + n^{-1/2} + D^{-1/2}$ up to the log order and the sparsity level with the finite J_1 and J_2 .

The following theorem provides the convergence rate of the predicted instantaneous volatility using the TIP-PCA method.

Theorem 3.1. *Suppose that Assumptions 3.1–3.5 hold, $J_1 = o(\sqrt{D})$, and $J_2 = o(\sqrt{n})$. As $D, n, m, J_1, J_2 \rightarrow \infty$, we have*

$$\begin{aligned} & \max_{j \leq n} |\tilde{c}_{D+1,j} - E[c_{D+1,j} | \mathcal{F}_D]| \\ &= O_P \left(\sqrt{\frac{\log D}{n}} + \frac{\varphi_D}{D} + \frac{J_2}{\sqrt{Dm^{1/4}}} + J_2 \left(\sqrt{\frac{\rho_m}{m^{1/8}}} + \rho_m \right) + J_2^{\frac{1}{2}-\frac{\kappa}{2}} + \frac{J_2}{\sqrt{n}} \right) \\ & \quad + O_P \left(\frac{J_1^{\frac{3}{2}}}{\sqrt{nm^{1/4}}} + J_1^{\frac{3}{2}} \left(\sqrt{\frac{\rho_m}{m^{1/8}}} + \rho_m \right) + J_1^{1-\frac{\kappa}{2}} + \frac{J_1^{\frac{3}{2}}}{\sqrt{D}} \right) \max_{l \leq J_1} \sup_x |\phi_l(x)|. \end{aligned}$$

Theorem 3.1 indicates that the proposed TIP-PCA consistently predicts the one-day-ahead instantaneous volatility. As discussed in Remark 3.2, we have $\sqrt{\frac{\rho_m}{m^{1/8}}}$, $D^{-1/2}$, $n^{-1/2}$, and J_i 's terms. To predict the one-day-ahead instantaneous volatility process, we need to learn the interday time series dynamics and intraday periodic patterns. Thus, the terms $D^{-1/2}$ and $n^{-1/2}$ are usual costs. We note that the difference from the result in Proposition 3.1 arises from the estimation of the nonparametric function $g_k(\cdot)$. Specifically, out-of-sample predictions with any covariates \mathbf{x} necessitate additional costs such as $J_1^{1/2}$ and the supremum term $\max_{l \leq J_1} \sup_x |\phi_l(x)|$, whereas the result in Proposition 3.1 does not require them because it is based on in-sample prediction.

4 Simulation Study

In this section, we conducted simulations to examine the finite sample performance of the proposed TIP-PCA method. We first generated high-frequency observations as follows: for $i = 1, \dots, D + 1$, $s = 0, \dots, m$, and $t_s = s/m$,

$$\begin{aligned} Y_{i,t_s} &= X_{i,t_s} + e_{i,t_s}, \\ dX_{i,t} &= (\mu - \sigma_{i,t}^2/2)dt + \sigma_{i,t}dB_{i,t} + J_{i,t}dP_{i,t}, \\ \sigma_{i,t_s}^2 &= \tilde{\sigma}_i^2 h(t_s) + \varepsilon_{i,t_s}, \end{aligned}$$

where we set microstructure noise as $e_{i,t_s} \sim \mathcal{N}(0, 0.0005^2)$ and the initial value as $X_{1,0} = 1$; $B_{i,t}$ is a standard Brownian motion; for the jump part, we set $J_{i,t} \sim \mathcal{N}(-0.01, 0.02^2)$ and $P_{i,t+\Delta} - P_{i,t} \sim \text{Poisson}(36\Delta/252)$; $\tilde{\sigma}_i = b_0 + b_1\tilde{\sigma}_{i-1} + b_2\frac{1}{5}\sum_{s=1}^5\tilde{\sigma}_{i-s} + b_3\frac{1}{22}\sum_{s=1}^{22}\tilde{\sigma}_{i-s} + \zeta_i$, $\zeta_i \sim \mathcal{N}(0, 1)$, $h(t_s) = \gamma_0 + \gamma_1(t_s - 0.6)^2$, and $\varepsilon_{i,t_s} = q(t_s)\xi_{i,t_s}$, where $q(t_s)^2 = 0.1 + 0.5(2t_s - 1)^2$ and $\xi_{i,t_s} \sim \mathcal{N}(0, 0.01^2)$. The model parameters were set to be

$$\begin{aligned} \mu &= 0.05/252, \gamma_0 = 0.04/252, \gamma_1 = 0.5/252, \\ b_0 &= 0.5, b_1 = 0.372, b_2 = 0.343, b_3 = 0.224. \end{aligned}$$

The normalized parameter values above imply the daily time unit, and we adapted the estimated coefficients studied in [Corsi \(2009\)](#) to generate $\tilde{\sigma}_i$. We note that in each simulation, the instantaneous volatility process was generated until all instantaneous volatility values were positive, based on the data-generating process described above. We set $m = 23,400$, which indicates that the data are observed every second over a period of 6.5 trading hours per day.

For each simulation, we used the jump robust pre-averaging method ([Figuroa-López and Wu, 2022](#)) to estimate the instantaneous volatility, $c_{i,\tau} = \sigma_{i,\frac{\tau}{n}}^2$, at a frequency of every 5 or

10 minutes (i.e., $n = m/300$ or $m/600$) for each i -th day as follows: for $\tau = 1, \dots, n$,

$$\widehat{c}_{i,\tau} = \frac{1}{\phi_{k_n}(g)} \sum_{s=1}^{m-k_n+1} K_{b_m}(t_{s-1} - \frac{\tau}{n}) \left(\bar{Y}_{i,s}^2 - \frac{1}{2} \widehat{Y}_{i,s} \right) \mathbf{1}_{\{|\bar{Y}_{i,s}| \leq \nu_m\}}, \quad (4.1)$$

where $K_b(x) = K(x/b)/b$, the bandwidth size $b_m = 1/n$, the weight function $g(x) = 2x \wedge (1-x)$,

$$\begin{aligned} \bar{Y}_{i,s} &= \sum_{l=1}^{k_m-1} g\left(\frac{l}{k_m}\right) (Y_{i,t_{s+l}} - Y_{i,t_{s+l-1}}), & \phi_{k_m}(g) &= \sum_{i=1}^{k_m} g\left(\frac{i}{k_m}\right)^2, \\ \widehat{Y}_{i,s} &= \sum_{l=1}^{k_m} \left\{ g\left(\frac{l}{k_m}\right) - g\left(\frac{l-1}{k_m}\right) \right\}^2 (Y_{i,t_{s+l}} - Y_{i,t_{s+l-1}})^2, \end{aligned}$$

$\mathbf{1}_{\{\cdot\}}$ is an indicator function, and $\nu_m = 1.8\sqrt{\text{BPV}}(k_m/m)^{0.47}$, where the bipower variation $\text{BPV} = \frac{\pi}{2} \sum_{s=2}^m |Y_{i,t_{s-1}} - Y_{i,t_{s-2}}| |Y_{i,t_s} - Y_{i,t_{s-1}}|$. We used the uniform kernel function and the data-driven approach to obtain the preaveraging window size, k_m , as suggested in Section 3.1 of [Figueroa-López and Wu \(2022\)](#).

With the instantaneous volatility estimates spanning D days, $\widehat{\Sigma}_{D,n} = (\widehat{c}_{i,\tau})_{D \times n}$, we examined the out-of-sample performance of estimating the one-day-ahead instantaneous volatility process. For comparison, the TIP-PCA, AVE, AR, SARIMA, HAR, HAR-D, XGBoost, and PC methods were employed to predict $c_{D+1,\tau}$, for $\tau = 1, \dots, n$. Specifically, for the TIP-PCA, we utilized the ex-post daily, weekly, and monthly realized volatilities and the intraday time sequence $\{\frac{\tau}{n}\}_{\tau=1}^n$ as covariates for \mathbf{X} and \mathbf{W} , respectively. In addition, the additive polynomial basis with $J_1 = 2$ and $J_2 = 3$ are used for the sieve basis of TIP-PCA, as discussed in Section 2.3. AVE represents estimates obtained by the column mean of $\widehat{\Sigma}_{D,n}$. AR and HAR represent predicted values obtained with the autoregressive model of order 1 and the HAR model, respectively, within each column of $\widehat{\Sigma}_{D,n}$. PC represents the last row of the estimated low-rank matrix using the best rank- r matrix approximation based on $\widehat{\Sigma}_{D,n}$. For TIP-PCA and PC, we used rank 1 as suggested by the eigenvalue ratio method ([Ahn and Horenstein, 2013](#)). We note that AR and HAR account for the time series dynamics

while AVE and PC cannot. However, AR and HAR have the overparameterization problem. PC can partially explain the periodic pattern using the rank-one right singular vector, while other competitors cannot explicitly account for the pattern due to the random noise ε_{i,t_j} . To address intraday periodic patterns in addition to daily autoregressive time series dynamics, we also employed the existing methods below. SARIMA represents n -step ahead predicted values obtained with the seasonal ARIMA(1,1,1) (Sheppard, 2010), with the length of the seasonal cycle as n . HAR-D represents the modified HAR model that includes the diurnal effect and previous intraday component in addition to the ex-post daily, weekly, and monthly realized volatilities (for details, see Zhang et al., 2024). Lastly, to account for the nonlinear impacts, we included XGBoost (Chen and Guestrin, 2016), a decision-tree-based ensemble algorithm. We utilized the same hyperparameters as those specified in Zhang et al. (2024). We note that SARIMA, HAR-D, and XGBoost were conducted after vectorizing $\widehat{\Sigma}_{D,n}$ as a time series vector. We generated high-frequency data with $m = 23,400$ for 200 consecutive days. We used the subsampled log prices of the last $D = 50, 100, 150,$ and 200 days. To check the performance of the instantaneous volatility, we calculated the mean squared prediction errors (MSPE) as follows:

$$\frac{1}{n} \sum_{\tau=1}^n (\widetilde{c}_{D+1,\tau} - c_{D+1,\tau})^2,$$

where $\widetilde{c}_{D+1,\tau}$ is one of the above one-day-ahead instantaneous volatility estimators. Then, we calculated the sample average of MSPEs over 500 simulations.

Figure 1 presents the average MSPEs of one-day-ahead intraday instantaneous volatility estimators with $D = \{50, 100, 150, 200\}$ and $n = \{39, 78\}$. We note that for each simulation, since we used the subsampled data, the target future volatility is the same for each different D . Figure 1 makes evident that the TIP-PCA method demonstrates the best performance. This is because TIP-PCA can accurately predict the future instantaneous volatility process by leveraging the features of the HAR model and the U-shaped intraday volatility and handling of the overparameterization problem. Additionally, the MSPEs of TIP-PCA tend to decrease as the number of daily observations and intraday instantaneous volatility estimators




Figure 1: $\text{MSPE} \times 10^9$ for the TIP-PCA, AVE, AR, SARIMA, HAR, HAR-D, XGBoost, and PC.

increases. This finding aligns with the theoretical results in Section 3. In contrast, the MSPEs of AVE and AR increase as the number of daily observations increases. This may be because they do not include the HAR model feature and consider old information deemed unhelpful. SARIMA and XGBoost do not perform well because they do not incorporate the interday HAR model feature. Furthermore, the HAR method does not perform well due to its inability to integrate the U-shaped intraday volatility feature and the overparameterization problem. By considering the previous intraday component as well as the diurnal effect, HAR-D improves performance compared to HAR. However, HAR-D underperforms TIP-PCA because it cannot fully incorporate the intraday periodic pattern. We note that we also considered a simulation study for the higher rank case ($r = 2$) in Section S.1 of the online supplement. The result is similar.

5 Empirical Study

In this section, we applied the proposed TIP-PCA method to an intraday instantaneous volatility prediction using real high-frequency trading data. We obtained intraday data of the

S&P 500 index ETF (SPY) and ETFs that represent the 11 Global Industrial Classification Standard (GICS) sector index funds (XLC, XLY, XLP, XLE, XLF, XLV, XLI, XLB, XLRE, XLK, and XLU) from July 2021 to June 2022 from the TAQ database in the Wharton Research Data Services (WRDS) system. We collected high-frequency data subsampled every 1 second and excluded days with early stock market closures during this period. This subsampling helps reduce the effect of irregular observation time points (Li and Linton, 2023). We used the log prices and employed the jump robust pre-averaging estimation procedure defined in Section 4 to estimate the instantaneous variance at a frequency of every 10 minutes. Then, we conducted the TIP-PCA, AVE, AR, SARIMA, HAR, HAR-D, XGBoost, and PC methods as described in Section 4 using the in-sample period data to predict the one-day-ahead instantaneous volatilities. In addition, we added TIP-PCA-S, which is a simplified version of the proposed method. Specifically, we fit the HAR model based on the row mean of $\widehat{\Sigma}_{D,n}$ to predict the one-day-ahead integrated volatility value, while the U-shape is fitted using the column mean of $\widehat{\Sigma}_{D,n}$. The predicted volatility process is then obtained by multiplying the predicted integrated volatility by the fitted intraday vector. We note that TIP-PCA-S basically generates a rank one matrix by averaging each dimension of the volatility matrix. We used the rolling window scheme, where the in-sample period was 63 days (i.e., one quarter). The number of out-of-sample predictions was 7,371 (i.e., every 10 minutes for 189 days).

To measure the performance of the predicted instantaneous volatility, we utilized the mean squared prediction errors (MSPE) and QLIKE (Patton, 2011) as follows:

$$\text{MSPE} = \frac{1}{nq} \sum_{i=1}^q \sum_{j=1}^n (\tilde{c}_{D+i,j} - \widehat{c}_{D+i,j})^2,$$

$$\text{QLIKE} = \frac{1}{nq} \sum_{i=1}^q \sum_{j=1}^n \left(\log \tilde{c}_{D+i,j} + \frac{\widehat{c}_{D+i,j}}{\tilde{c}_{D+i,j}} \right),$$

where $\tilde{c}_{D+i,j}$ is one of the TIP-PCA, AVE, AR, SARIMA, HAR, HAR-D, XGBoost, and PC estimates. We predicted one-day-ahead conditional expected instantaneous volatilities using

Table 1: MSPEs and QLIKEs for the TIP-PCA, AVE, AR, SARIMA, HAR, HAR-D, XGBoost, PC, and TIP-PCA-S.

	TIP-PCA	AVE	AR	SARIMA	HAR	HAR-D	XGBoost	PC	TIP-PCA-S
MSPE $\times 10^9$									
SPY	2.239	3.213	2.903	3.082	3.328	2.942	2.843	2.407	2.433
XLC	3.308	4.491	4.055	4.200	4.794	4.719	4.105	3.455	3.782
XLY	8.606	11.992	10.327	10.375	12.162	12.399	12.533	9.157	9.507
XLP	0.593	0.798	0.713	0.828	0.804	0.752	0.719	0.604	0.627
XLE	8.471	11.792	9.516	9.399	10.017	10.105	10.919	8.559	8.909
XLF	2.380	3.510	2.854	3.121	3.085	2.897	3.118	2.524	2.507
XLV	0.918	1.223	1.069	1.150	1.214	1.172	1.334	0.953	1.052
XLI	1.817	2.568	2.344	2.325	2.792	2.373	2.313	1.851	1.986
XLB	2.182	2.993	2.605	2.819	2.965	2.891	2.912	2.254	2.376
XLRE	1.803	2.339	2.120	2.583	2.400	2.218	2.327	1.819	1.881
XLK	7.477	10.233	9.019	9.157	10.628	9.819	9.538	7.879	8.159
XLU	0.909	1.143	1.026	1.129	1.120	1.081	1.177	0.926	0.939
QLIKE									
SPY	-9.054	-8.932	-9.095	-6.983	-8.695	-7.098	-9.098	-9.171	-9.150
XLC	-8.871	-8.801	-8.898	-7.061	-1.790	-8.378	-8.897	-8.973	-8.970
XLY	-8.367	-8.179	-8.297	-7.600	-7.987	-6.081	-8.286	-8.379	-8.357
XLP	-9.477	-9.378	-9.444	-7.934	-9.316	-8.932	-9.450	-9.484	-9.473
XLE	-8.090	-8.013	-8.060	-7.993	-7.947	-7.971	-8.068	-8.093	-8.085
XLF	-8.393	-8.327	-8.379	-8.303	-8.325	-8.057	-8.377	-8.391	-8.392
XLV	-9.417	-9.302	-9.391	-6.121	-9.153	-6.680	-9.371	-9.451	-9.425
XLI	-9.134	-8.997	-9.109	-6.366	-8.579	-8.638	-9.122	-9.197	-9.358
XLB	-9.104	-8.947	-9.025	-6.564	-8.670	-8.378	-9.047	-9.118	-9.093
XLRE	-9.082	-8.986	-9.036	-7.208	-8.845	-8.244	-9.029	-9.075	-9.067
XLK	-8.487	-8.316	-8.452	-6.959	-8.172	-7.382	-8.442	-8.511	-8.500
XLU	-9.122	-9.057	-9.097	-7.974	-9.017	-8.990	-9.100	-9.117	-9.119

in-sample period data. Additionally, since we do not know the true conditional expected instantaneous volatility, to assess the significance of differences in prediction performances, we conducted the Diebold and Mariano (DM) test (Diebold and Mariano, 2002) based on MSPE and QLIKE. We compared the proposed TIP-PCA method with other methods. Importantly, since we conducted each hypothesis test multiple times repeatedly throughout this section, it is essential to control the False Discovery Rate (FDR). Hence, we adjusted the p -values using the Benjamini–Hochberg (BH) procedure (Benjamini and Hochberg, 1995) at a level $\alpha = 0.05$ to control the FDR across all hypothesis tests employed in this section.

Table 1 reports the results of MSPEs and QLIKEs, and Table 2 shows the adjusted p -

Table 2: The adjusted p -values for the DM test statistic based on MSPE and QLIKE for AVE, AR, SARIMA, HAR, HAR-D, XGBoost, PC, and TIP-PCA-S with respect to the TIP-PCA.

	AVE	AR	SARIMA	HAR	HAR-D	XGBoost	PC	TIP-PCA-S
MSPE								
SPY	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***
XLC	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***	0.013**	0.000***
XLY	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***
XLP	0.000***	0.001***	0.000***	0.000***	0.000***	0.000***	0.495	0.000***
XLE	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***	0.857	0.000***
XLF	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***
XLV	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***	0.141	0.000***
XLI	0.000***	0.001***	0.000***	0.000***	0.000***	0.000***	0.140	0.000***
XLB	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***	0.043**	0.000***
XLRE	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***	0.623	0.000***
XLK	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***
XLU	0.000***	0.002***	0.000***	0.000***	0.000***	0.000***	0.252	0.000***
QLIKE								
SPY	0.000***	0.084*	0.000***	0.003***	0.017**	0.077*	0.000***	0.000***
XLC	0.043**	0.379	0.000***	0.400	0.000***	0.312	0.000***	0.000***
XLY	0.000***	0.000***	0.000***	0.000***	0.001***	0.000***	0.000***	0.000***
XLP	0.000***	0.000***	0.001***	0.005***	0.001***	0.000***	0.000***	0.000***
XLE	0.000***	0.000***	0.000***	0.000***	0.141	0.000***	0.176	0.000***
XLF	0.000***	0.000***	0.000***	0.000***	0.004***	0.000***	0.038**	0.027**
XLV	0.000***	0.002***	0.000***	0.000***	0.101	0.000***	0.000***	0.011**
XLI	0.000***	0.046**	0.000***	0.007***	0.000***	0.276	0.000***	0.000***
XLB	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***
XLRE	0.000***	0.000***	0.000***	0.000***	0.011**	0.000***	0.094*	0.000***
XLK	0.000***	0.000***	0.001***	0.000***	0.000***	0.014**	0.000***	0.000***
XLU	0.000***	0.000***	0.000***	0.000***	0.000***	0.000***	0.005***	0.000***

Note: The adjusted p -values are based on the Benjamini–Hochberg (BH) procedure (Benjamini and Hochberg, 1995) at a level $\alpha = 0.05$. ***, **, and * indicate rejection of the null hypothesis at significance levels of 1%, 5%, and 10%, respectively.

values for the DM tests. From Tables 1 and 2, we find that the TIP-PCA method exhibits the best performance overall. This may be because the projection method, utilizing covariates such as ex-post realized volatility information and the U-shaped intraday volatility feature, contributes to enhancing the accuracy of instantaneous volatility predictions. We note that PC also performs well compared to other competing methods. This may be because removing the noise component using the PC method is crucial. Additionally, the dynamics of previous intraday volatility significantly influence one-day-ahead intraday volatility fore-

Table 3: Number of cases where the adjusted p -value is greater than 0.05 for TIP-PCA, AVE, AR, SARIMA, HAR, HAR-D, XGBoost, PC, and TIP-PCA-S across 12 ETFs at each $q_0 = \{0.01, 0.02, 0.05, 0.1, 0.2\}$ based on the LRuc, LRcc, and DQ tests.

q_0	LRuc					LRcc					DQ				
	0.01	0.02	0.05	0.1	0.2	0.01	0.02	0.05	0.1	0.2	0.01	0.02	0.05	0.1	0.2
TIP-PCA	9	12	12	12	12	10	12	12	12	12	6	6	7	8	10
AVE	3	3	6	12	12	3	4	5	12	12	3	3	4	4	7
AR	4	5	11	12	12	7	7	11	12	12	4	4	5	7	9
SARIMA	1	1	1	3	10	7	9	11	12	12	2	2	2	3	6
HAR	1	2	4	12	12	0	1	4	11	12	3	4	4	5	7
HAR-D	1	0	4	7	12	7	10	12	12	12	2	2	3	4	6
XGBoost	3	4	8	11	12	7	10	11	12	12	3	3	5	6	9
PC	6	8	12	12	12	7	9	12	12	12	5	5	7	8	10
TIP-PCA-S	4	6	11	12	12	8	10	12	12	12	4	5	6	8	10

Note: The adjusted p -values are based on the Benjamini–Hochberg (BH) procedure (Benjamini and Hochberg, 1995) at a level $\alpha = 0.05$.

casts, which indicates a strong AR structure in the interday dimension. TIP-PCA-S is also comparable to the proposed TIP-PCA in some cases. This may be because both TIP-PCA and TIP-PCA-S utilize the same information using the HAR model structure within the instantaneous volatility matrix. To further evaluate the performance of the proposed method, we implemented one-day-ahead Value at Risk (VaR) estimation as described below.

We also evaluated the performance of the proposed method in estimating one-day-ahead 10-minute frequency Value at Risk (VaR). In particular, we first predicted the one-day-ahead conditional expected instantaneous volatilities using the TIP-PCA, AVE, AR, SARIMA, HAR, HAR-D, XGBoost, and PC procedures using the in-sample period data. We then calculated the quantiles using historical standardized 10-minute returns. Specifically, we standardized in-sample 10-minute returns using estimated conditional instantaneous volatilities. We then derived sample quantiles for 0.01, 0.02, 0.05, 0.1, and 0.2. Using the sample quantile estimates and predicted instantaneous volatility, we obtained the one-day-ahead 10-minute frequency VaR values for each prediction method. We used a fixed in-sample period as one quarter and implemented a rolling window scheme. The out-of-sample period was considered to be from October 2021 to June 2022.

To backtest the estimated VaR, we conducted the likelihood ratio unconditional coverage (LRuc) test (Kupiec, 1995), the likelihood ratio conditional coverage (LRcc) test (Christoffersen, 1998), and the dynamic quantile (DQ) test with lag 4 (Engle and Manganelli, 2004). Table 3 reports the number of cases where the adjusted p -value using the BH procedure is greater than 0.05 for the 12 ETFs at each $q_0 = \{0.01, 0.02, 0.05, 0.1, 0.2\}$ quantile, based on the LRuc, LRcc, and DQ tests. In Table 3, we find that the TIP-PCA method consistently outperforms in all hypothesis tests. We note that both PC and TIP-PCA-S also perform well compared to other methods. However, the number of cases for TIP-PCA exceeds those of PC and TIP-PCA-S, particularly for the lower quantiles, which are difficult to predict. This is because TIP-PCA-S cannot filter out the noise component using the given singular vector structure, and PC cannot explain the complex dynamic structure. In contrast, TIP-PCA can filter out the noise component through the projection approach and catch the complex dynamics using the inter-day HAR and intra-day U-shape structures. This outcome confirms that the proposed TIP-PCA method significantly contributes to the improved prediction accuracy of future instantaneous volatilities and enhanced risk management.

6 Conclusion

This paper introduces a novel intraday instantaneous volatility prediction procedure. The proposed Two-sIde-Projected-PCA (TIP-PCA) method leverages both interday and intraday volatility dynamics based on the semiparametric structure of the low-rank matrix of the instantaneous volatility process. We establish the asymptotic properties of TIP-PCA and its future instantaneous volatility estimators. In the empirical study, concerning the out-of-sample performance of predicting the one-day-ahead instantaneous volatility process, TIP-PCA outperforms other conventional methods. This finding confirms that both the HAR model structure on interday dynamics and the U-shaped pattern on intraday dynamics contribute to predicting the future instantaneous volatility process.

In this paper, we focus on the instantaneous volatility process for a single asset. In practice, we often need to handle a large number of assets. Thus, it is important and interesting to extend the study to predict the instantaneous volatility process of many assets. However, to do this, cross-sectionally, we encounter another curse of dimensionality. Therefore, technically, it is a demanding task to handle both the cross-sectional curse of dimensionality and the intraday curse of dimensionality. We leave this for a future study.

References

- ADMATI, A. R. AND P. PFLEIDERER (1988): “A theory of intraday patterns: Volume and price variability,” *The Review of Financial Studies*, 1, 3–40.
- AHN, S. C. AND A. R. HORENSTEIN (2013): “Eigenvalue ratio test for the number of factors,” *Econometrica*, 81, 1203–1227.
- AÏT-SAHALIA, Y., J. FAN, AND D. XIU (2010): “High-frequency covariance estimates with noisy and asynchronous financial data,” *Journal of the American Statistical Association*, 105, 1504–1517.
- AÏT-SAHALIA, Y. AND J. YU (2009): “High frequency market microstructure noise estimates and liquidity measures,” *Annals of Applied Statistics*, 3, 422–457.
- ANDERSEN, T. G. AND T. BOLLERSLEV (1997): “Intraday periodicity and volatility persistence in financial markets,” *Journal of Empirical Finance*, 4, 115–158.
- ANDERSEN, T. G., T. BOLLERSLEV, F. X. DIEBOLD, AND P. LABYS (2003): “Modeling and forecasting realized volatility,” *Econometrica*, 71, 579–625.
- ANDERSEN, T. G., M. THYRSGAARD, AND V. TODOROV (2019): “Time-varying periodicity in intraday volatility,” *Journal of the American Statistical Association*, 114, 1695–1707.

- BAI, J. AND S. NG (2002): “Determining the number of factors in approximate factor models,” *Econometrica*, 70, 191–221.
- BARNDORFF-NIELSEN, O. E., P. R. HANSEN, A. LUNDE, AND N. SHEPHARD (2008): “Designing realized kernels to measure the ex post variation of equity prices in the presence of noise,” *Econometrica*, 76, 1481–1536.
- (2011): “Multivariate realised kernels: consistent positive semi-definite estimators of the covariation of equity prices with noise and non-synchronous trading,” *Journal of Econometrics*, 162, 149–169.
- BENJAMINI, Y. AND Y. HOCHBERG (1995): “Controlling the false discovery rate: a practical and powerful approach to multiple testing,” *Journal of the Royal statistical society: series B (Methodological)*, 57, 289–300.
- BIBINGER, M., N. HAUTSCH, P. MALEC, AND M. REISS (2014): “Estimating the quadratic covariation matrix from noisy observations: Local method of moments and efficiency,” *The Annals of Statistics*, 42, 1312–1346.
- BOIVIN, J. AND S. NG (2006): “Are more data always better for factor analysis?” *Journal of Econometrics*, 132, 169–194.
- CANDES, E. J. AND Y. PLAN (2010): “Matrix completion with noise,” *Proceedings of the IEEE*, 98, 925–936.
- CHEN, E. Y., D. XIA, C. CAI, AND J. FAN (2020): “Semiparametric tensor factor analysis by iteratively projected SVD,” *arXiv preprint arXiv:2007.02404*.
- CHEN, T. AND C. GUESTRIN (2016): “Xgboost: A scalable tree boosting system,” in *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, 785–794.

- CHEN, X. (2007): “Large sample sieve estimation of semi-nonparametric models,” *Handbook of Econometrics*, 6, 5549–5632.
- CHO, J., D. KIM, AND K. ROHE (2017): “Asymptotic theory for estimating the singular vectors and values of a partially-observed low rank matrix with noise,” *Statistica Sinica*, 1921–1948.
- CHRISTENSEN, K., S. KINNEBROCK, AND M. PODOLSKIJ (2010): “Pre-averaging estimators of the ex-post covariance matrix in noisy diffusion models with non-synchronous data,” *Journal of Econometrics*, 159, 116–133.
- CHRISTOFFERSEN, P. F. (1998): “Evaluating interval forecasts,” *International Economic Review*, 841–862.
- CORSI, F. (2009): “A simple approximate long-memory model of realized volatility,” *Journal of Financial Econometrics*, 7, 174–196.
- DIEBOLD, F. X. AND R. S. MARIANO (2002): “Comparing predictive accuracy,” *Journal of Business & Economic Statistics*, 20, 134–144.
- ENGLE, R. F. AND S. MANGANELLI (2004): “CAViaR: Conditional autoregressive value at risk by regression quantiles,” *Journal of Business & Economic Statistics*, 22, 367–381.
- ENGLE, R. F. AND M. E. SOKALSKA (2012): “Forecasting intraday volatility in the US equity market. Multiplicative component GARCH,” *Journal of Financial Econometrics*, 10, 54–83.
- FAN, J., A. FURGER, AND D. XIU (2016a): “Incorporating global industrial classification standard into portfolio allocation: A simple factor-based large covariance matrix estimator with high-frequency data,” *Journal of Business & Economic Statistics*, 34, 489–503.
- FAN, J. AND D. KIM (2018): “Robust high-dimensional volatility matrix estimation for

- high-frequency factor model,” *Journal of the American Statistical Association*, 113, 1268–1283.
- FAN, J., Y. LIAO, AND W. WANG (2016b): “Projected principal component analysis in factor models,” *The Annals of Statistics*, 44, 219.
- FAN, J., H. LIU, AND W. WANG (2018a): “Large covariance estimation through elliptical factor models,” *Annals of Statistics*, 46, 1383.
- FAN, J., W. WANG, AND Y. ZHONG (2018b): “An ∞ eigenvector perturbation bound and its application to robust covariance estimation,” *Journal of Machine Learning Research*, 18, 1–42.
- FAN, J. AND Y. WANG (2007): “Multi-scale jump and volatility analysis for high-frequency financial data,” *Journal of the American Statistical Association*, 102, 1349–1362.
- (2008): “Spot volatility estimation for high-frequency data,” *Statistics and its Interface*, 1, 279–288.
- FIGUEROA-LÓPEZ, J. E. AND B. WU (2022): “Kernel estimation of spot volatility with microstructure noise using pre-averaging,” *Econometric Theory*, 1–50.
- FOSTER, D. P. AND D. B. NELSON (1996): “Continuous record asymptotics for rolling sample variance estimators,” *Econometrica*, 64, 139–174.
- HANSEN, P. R., Z. HUANG, AND H. H. SHEK (2012): “Realized GARCH: a joint model for returns and realized measures of volatility,” *Journal of Applied Econometrics*, 27, 877–906.
- HONG, H. AND J. WANG (2000): “Trading and returns under periodic market closures,” *The Journal of Finance*, 55, 297–354.
- JACOD, J., Y. LI, P. A. MYKLAND, M. PODOLSKIJ, AND M. VETTER (2009): “Microstructure noise in the continuous case: the pre-averaging approach,” *Stochastic Processes and their Applications*, 119, 2249–2276.

- KIM, D. AND J. FAN (2019): “Factor GARCH-Itô models for high-frequency data with application to large volatility matrix prediction,” *Journal of Econometrics*, 208, 395–417.
- KIM, D. AND Y. WANG (2016): “Unified discrete-time and continuous-time models and statistical inferences for merged low-frequency and high-frequency financial data,” *Journal of Econometrics*, 194, 220–230.
- KIM, D., Y. WANG, AND J. ZOU (2016): “Asymptotic theory for large volatility matrix estimation based on high-frequency financial data,” *Stochastic Processes and their Applications*, 126, 3527–3577.
- KRISTENSEN, D. (2010): “Nonparametric filtering of the realized spot volatility: A kernel-based approach,” *Econometric Theory*, 26, 60–93.
- KUPIEC, P. H. (1995): “Techniques for Verifying the Accuracy of Risk Measurement Models,” *The Journal of Derivatives*, 3, 73–84.
- LI, Z. M. AND O. LINTON (2023): “Robust estimation of integrated and spot volatility,” *Journal of Econometrics*, 105614.
- MANCINI, C., V. MATTIUSI, AND R. RENÒ (2015): “Spot volatility estimation using delta sequences,” *Finance and Stochastics*, 19, 261–293.
- ONATSKI, A. (2010): “Determining the number of factors from empirical distribution of eigenvalues,” *The Review of Economics and Statistics*, 92, 1004–1016.
- PATTON, A. J. (2011): “Volatility forecast comparison using imperfect volatility proxies,” *Journal of Econometrics*, 160, 246–256.
- SHEPHARD, N. AND K. SHEPPARD (2010): “Realising the future: forecasting with high-frequency-based volatility (HEAVY) models,” *Journal of Applied Econometrics*, 25, 197–231.

- SHEPPARD, K. (2010): “Financial Econometrics Notes,” *University of Oxford*, 333–426.
- SHIN, M., D. KIM, AND J. FAN (2023): “Adaptive robust large volatility matrix estimation based on high-frequency financial data,” *Journal of Econometrics*, 237, 105514.
- SONG, X., D. KIM, H. YUAN, X. CUI, Z. LU, Y. ZHOU, AND Y. WANG (2021): “Volatility analysis with realized GARCH-Itô models,” *Journal of Econometrics*, 222, 393–410.
- TODOROV, V. (2019): “Nonparametric spot volatility from options,” *The Annals of Applied Probability*, 29, 3590–3636.
- TODOROV, V. AND Y. ZHANG (2023): “Bias reduction in spot volatility estimation from options,” *Journal of Econometrics*, 234, 53–81.
- VERSHYNIN, R. (2010): “Introduction to the non-asymptotic analysis of random matrices,” *arXiv preprint arXiv:1011.3027*.
- WANG, W. AND J. FAN (2017): “Asymptotics of empirical eigenstructure for high dimensional spiked covariance,” *The Annals of Statistics*, 45, 1342.
- XIU, D. (2010): “Quasi-maximum likelihood estimation of volatility with high frequency data,” *Journal of Econometrics*, 159, 235–250.
- ZHANG, C., Y. ZHANG, M. CUCURINGU, AND Z. QIAN (2024): “Volatility forecasting with machine learning and intraday commonality,” *Journal of Financial Econometrics*, 22, 492–530.
- ZHANG, L. (2006): “Efficient estimation of stochastic volatility using noisy observations: A multi-scale approach,” *Bernoulli*, 12, 1019–1043.
- (2011): “Estimating covariation: Epps effect, microstructure noise,” *Journal of Econometrics*, 160, 33–47.

ZHANG, L., P. A. MYKLAND, AND Y. AÏT-SAHALIA (2005): “A tale of two time scales: Determining integrated volatility with noisy high-frequency data,” *Journal of the American Statistical Association*, 100, 1394–1411.

ZU, Y. AND H. P. BOSWIJK (2014): “Estimating spot volatility with high-frequency financial data,” *Journal of Econometrics*, 181, 117–135.