

# Forecasting intraday financial time series with sieve bootstrapping and dynamic updating

Han Lin Shang<sup>1</sup>  | Kaiying Ji<sup>2</sup>

<sup>1</sup>Department of Actuarial Studies and Business Analytics, Macquarie University, Sydney, New South Wales, Australia

<sup>2</sup>Discipline of Accounting, Governance and Regulation, The University of Sydney, Sydney, New South Wales, Australia

## Correspondence

Han Lin Shang, Department of Actuarial Studies and Business Analytics, Macquarie University, 4 Eastern Road, Sydney, NSW 2109, Australia.  
Email: [hanlin.shang@mq.edu.au](mailto:hanlin.shang@mq.edu.au)

## Funding information

Australian Research Council Discovery Project, Grant/Award Number: ARC DP230102250

## Abstract

Intraday financial data often take the form of a collection of curves that can be observed sequentially over time, such as intraday stock price curves. These curves can be viewed as a time series of functions observed on equally spaced and dense grids. Due to the curse of dimensionality, high-dimensional data pose challenges from a statistical aspect; however, it also provides opportunities to analyze a rich source of information so that the dynamic changes within short-time intervals can be better understood. We consider a sieve bootstrap method to construct 1-day-ahead point and interval forecasts in a model-free way. As we sequentially observe new data, we also implement two dynamic updating methods to update point and interval forecasts for achieving improved accuracy. The forecasting methods are validated through an empirical study of 5-min cumulative intraday returns of the S&P/ASX All Ordinaries Index.

## KEYWORDS

function-on-function linear regression, functional principal component analysis, high-frequency financial data, penalized least squares, vector autoregressive model

## 1 | INTRODUCTION

High-frequency financial data are observations on financial instruments taken at a finer time scale, such as time-stamped transaction-by-transaction or tick-by-tick data. Advances in computer recording and storage facilitate the presence of such data, which is often used to study market microstructure related to price discovery and market efficiency (see, e.g., Engle, 2000). Several pioneer articles on the application of high-frequency financial data sets are provided by Andersen (2000), Campbell et al. (1997), Dacarogna et al. (2001), Ghysels (2000), Goodhart and O'Hara (1997), Gouriéroux and Jasiak (2001), Lyons (2001), Wood (2000), Tsay (2010), and Peña and Tsay (2021). This study extends the existing literature

by presenting a functional time-series method with sieve bootstrapping to obtain well-calibrated interval forecasts. It applies the methodology to the S&P/ASX All Ordinaries Index (ASX code: XAO).

High-frequency financial data possess unique opportunities absent in low-frequency data. Traditional financial models, such as the celebrated capital asset pricing model (Lintner, 1965; Sharpe, 1964), consider stock return using low-frequency spot prices, where the price changes between the two spots are largely ignored. If prices are modeled as a univariate time series of discrete observations, the underlying process that generates these observations cannot be fully discovered. Suppose the closing prices are the same for two consecutive days, indicating no movement in the stock market. In contrast, the

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2023 The Authors. *Journal of Forecasting* published by John Wiley & Sons Ltd.

high-frequency intraday information is particularly interesting in this study because it contains finer trading information and allows us to investigate the movement in an underlying stock. One can study the patterns hidden in the intraday price curves and how the shapes of the intraday price curves change daily.

A high-frequency intraday financial data set is an example of (dense) functional data, represented in a graphical form of curves. Functional time series arise in many fields, including finance (see, e.g., Kokoszka & Zhang, 2012; Kokoszka et al., 2017; Wang et al., 2008). Functional time series can be of two types: On one hand, functional time series can arise by separating a continuous time record into natural (subjective) consecutive intervals. On the other hand, functional time series can also arise when observations in a period are considered as finite realizations of a continuous function, such as maturity-specific yield curves (see, e.g., Hays et al., 2012). In either case, the object of interest is a discrete time series of functions with a continuum. The continuum can either be time (e.g., intraday) or other continuous variables (e.g., maturity). In Section 2, we describe our data set—5-min cumulative intraday returns (CIDRs) of the XAO. This data set belongs to the first type of functional time series.

The advantages of functional time series include (1) we can study the temporal correlation of an intraday functional object and learn about how the correlation progress over days; (2) we can handle missing values via interpolation or smoothing techniques; (3) although not the focus of this paper, we can study not only the level but also the derivatives of functions as a means of reflecting the speed of change (see Hooker & Shang, 2022; Shang, 2019); (4) by converting a univariate time series in a real-valued space to a time series of functions in function space, we overcome the “curse of dimensionality” due to natural ordering.

To model a time series of functions, parametric and nonparametric techniques exist in the statistical literature. In the parametric tools, Bosq (2000) proposed a functional autoregressive of order 1 (FAR(1)) and derived one-step-ahead forecasts based on a regularized form of Yule–Walker equations. The FAR(1) model was later extended to a FAR( $p$ ) model by Kokoszka and Reimherr (2013), where the order  $p$  could be determined using a sequential testing procedure. Klepsch and Klüppelberg (2017) proposed a functional moving average process and introduced an innovations algorithm. By combining autoregressive and moving average, Klepsch et al. (2017) proposed a functional autoregressive and moving average model. While these techniques are designed for analyzing short-memory stationary functional time series, Li et al. (2020, 2021, 2023) considered long-range-dependent

functional time series and proposed a functional autoregressive fractionally integrated moving average model.

In the nonparametric tools, functional principal component analysis has been studied extensively to summarize the characteristics of an infinite-dimensional object by a set of finite-dimensional bases. The functional principal component analysis decomposes a time series of functions into a set of functional principal components and their associated principal component scores. The temporal dependence in a functional time series is inherited by the correlation among principal component scores. We adopt the approach of Aue et al. (2015), who applied a multivariate time-series forecasting method, such as a vector autoregressive model, to model and forecast multiple sets of estimated principal component scores jointly. Point forecasts can be obtained based on the historical functions and estimated functional principal components. Because this method uses a time-series forecasting method, it is referred to as the “TS method” in Section 3.1.

To quantify forecast uncertainty, a sieve bootstrap procedure of Paparoditis and Shang (2023) is implemented for constructing pointwise prediction intervals and uniform prediction bands. As Box (1976) stated “all models are wrong, but some are useful.” Under this principle, point forecasts obtained from a statistical time-series model are a proxy of future expectations and are of little use for purposes such as risk management. The sieve bootstrap in Section 3.2 takes account of model misspecification error as an additional source of uncertainty, thus improving calibration. The success of this sieve bootstrap is constructing a prediction error distribution that includes the model misspecification error and thus achieves improved calibration.

Forecasting the expected values of stock returns is a questionable exercise. According to Fama's (1970) weak form of efficient markets hypothesis, forecasting stock returns based on historical data is impossible. However, several studies have shown that forecasting intraday returns are possible. In a cross-section of stocks, Heston et al. (2010) demonstrated a pattern of return continuation at half-hour intervals that were multiples of a trading day, and this pattern can last for at least 40 trading days. Gao et al. (2017) demonstrated that first half-hour returns on the market predict the last half-hour returns. This study extends the findings of Gao et al. (2017) and demonstrates that dynamic updating methods can be applied throughout the day with various degrees of accuracy and can be generalized to other capital markets, such as Australia.

When partially observed intraday data in the most recent curve becomes available, we should incorporate this new information to improve forecast accuracy. We

present two dynamic updating methods, namely, penalized least squares (PLS) and function-on-function linear regression, for updating point and interval forecasts in Section 4. The dynamic updating methods can incorporate any intraday price jumps and obtain revised point and interval forecasts. By inputting sieve bootstrap forecasts into the dynamic updating methods, we may obtain improved forecast accuracy, which is one of the primary contributions of the paper. Using the forecast error measures in Section 5, via the CIDRs of the XAO, we evaluate and compare point and interval forecast accuracies in Section 6. Conclusions are given in Section 7, along with some ideas on how the methodology can be further extended.

## 2 | CIDRS OF S&P/ASX ALL ORDINARIES (XAO)

The All Ordinaries is designed to measure the 500 largest companies in the Australian equities market, drawn from eligible companies listed on the Australian Securities Exchange. According to the Union Bank of Switzerland, the Australian market is highly concentrated, with large financial, resources, and technology companies. As the first major financial market to open each day, the Australian Securities Exchange is a world leader in raising capital and provides active managers with opportunities to beat the market. We focus on XAO to add some fresh evidence to the existing forecasting literature, which largely overlooks such a major capital market. The 5-min intraday close prices of XAO from January 4 to December 23, 2021, were downloaded from Refinitiv (<https://select.datascope.refinitiv.com/DataScope/>). We consider 1 year of data to avoid any structural change. Following the work of Hansen and Lunde (2006), we also consider a 5-min sampling frequency which is a popular choice when avoiding the adverse effects of microstructure noise.

Let  $P_t(u_i)$ ,  $t = 1, 2, \dots, n$ ,  $i = 2, 3, \dots, \tau$ , and  $\tau = 75$  be the 5-min close price of the XAO at intraday time  $u_i$  between 10:00 and 16:10 Sydney time on day  $t$ . The close time of 16:10 is to allow all trades that take place on the close transaction at the same price regardless of what price an investor bids or offers. Any overnight trading will be reflected at the beginning close price the next day. For a stationary transformation, we compute CIDRs as discussed in Kokoszka et al. (2017), defined as

$$\mathcal{X}_t(u_i) = 100 \times [\ln P_t(u_i) - \ln P_t(u_1)], \quad i = 2, 3, \dots, \tau, \quad (2.1)$$

where  $\ln(\cdot)$  represents the natural logarithm. There are other versions of intraday return curves (see Kokoszka

et al., 2017; Rice et al., 2020). We choose the CIDR curve to eliminate the possible jump at the first close price. The CIDR curves preserve intraday movements, enhancing our understanding of the evolutions of intraday returns for XAO. From (2.1), we take the inverse transformation to obtain the 5-min intraday price

$$P_t(u_i) = \exp^{\frac{\mathcal{X}_t(u_i)}{100}} \times P_t(u_1).$$

Because  $\tau > \lfloor \sqrt{n} \rfloor$ , we have a dense functional data setting (Zhang & Wang, 2016). Via a linear interpolation, continuous curves can be constructed from these 74 CIDRs, and we form a time series of functions  $\mathcal{X}_t(u)$ . In the XAO data set, there are 74 five-minute CIDRs, representing around 6 h of trading at the Australian Securities Exchange on a trading day. A univariate time series of 18,500 = 74 × 250 discrete returns was converted into  $n = 250$  days of CIDRs. Figure 1a presents a univariate time-series display of CIDRs, whereas Figure 1b displays the same data as a time series of functions. From Figure 1a, the data series appears to have a constant mean, persistence, and no apparent volatility clustering. Through a functional KPSS test of Horváth et al. (2014), the series is trend stationary with a  $p$ -value of 0.737, implying no presence of a unit root.

## 3 | FUNCTIONAL TIME-SERIES FORECASTING METHOD

### 3.1 | Functional principal component regression

Let  $(\mathcal{X}_t, t \in \mathbb{Z})$  be an arbitrary stationary functional time series. Each function is a square-integrable function in Hilbert space  $\mathcal{H}$  satisfying  $\|\mathcal{X}_t\|^2 = \int_{\mathcal{I}} \mathcal{X}_t^2(u) du < \infty$  where  $\mathcal{I} \in [u_2, u_\tau]$  denotes a function support. While  $u_2$  symbolizes the first value of CIDRs recorded at 10:05, and  $u_\tau$  denotes the last value of CIDRs recorded at 16:10 on a trading day. For a time series of functions  $[\mathcal{X}_1(u), \mathcal{X}_2(u), \dots, \mathcal{X}_n(u)]$ , the mean function can be estimated by  $\bar{\mathcal{X}}(u) = \frac{1}{n} \sum_{t=1}^n \mathcal{X}_t(u)$ . The covariance function can be estimated by

$$\text{cov}[\mathcal{X}(u), \mathcal{X}(v)] = \text{E}\{[\mathcal{X}(u) - \bar{\mathcal{X}}(u)][\mathcal{X}(v) - \bar{\mathcal{X}}(v)]\}.$$

Via Mercer's lemma, the covariance function can be approximated by orthonormal eigenfunctions,

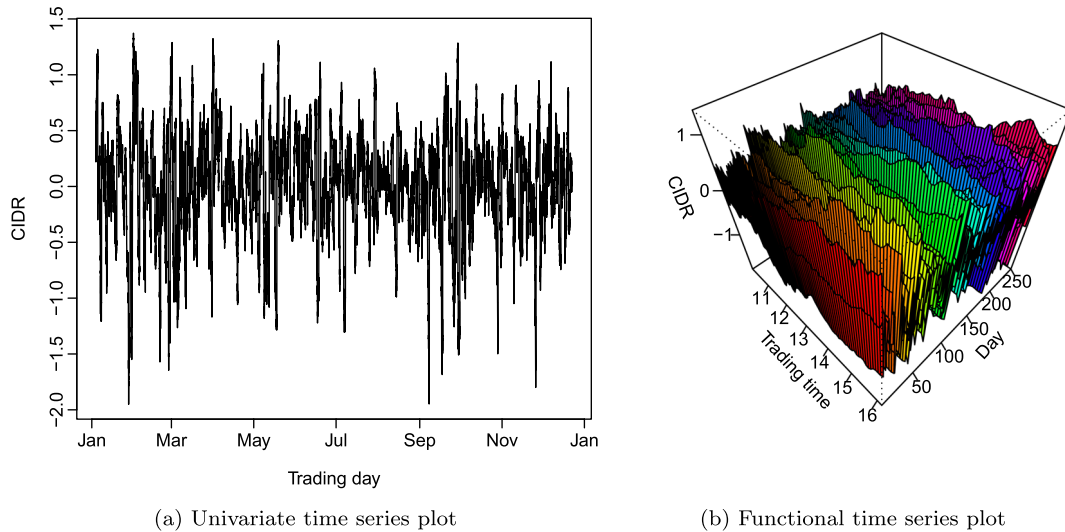


FIGURE 1 Graphical displays of CIDRs of the XAO from January 4 to December 23, 2021.

$$\text{cov}[\mathcal{X}(u), \mathcal{X}(v)] = \sum_{k=1}^{\infty} \hat{\lambda}_k \hat{\phi}_k(u) \hat{\phi}_k(v), \quad u, v \in \mathcal{I},$$

$$K = \underset{1 \leq k \leq k_{\max}}{\text{argmin}} \left\{ \frac{\hat{\lambda}_{k+1}}{\hat{\lambda}_k} \times \mathbb{1} \left( \frac{\hat{\lambda}_k}{\hat{\lambda}_1} \geq \nu \right) + \mathbb{1} \left( \frac{\hat{\lambda}_k}{\hat{\lambda}_1} < \nu \right) \right\}, \quad (3.3)$$

where  $\hat{\phi}_k(u)$  denotes the  $k$ th orthonormal functional principal component,  $\hat{\lambda}_k$  denotes the  $k$ th eigenvalue, and  $\hat{\lambda}_1 \geq \hat{\lambda}_2 \dots > 0$ .

Orthonormal eigenfunctions can well approximate realizations of a stochastic process. Via the Karhunen–Loève expansion, the functional realization  $\mathcal{X}_t(u)$  can be expressed as

$$\mathcal{X}_t(u) = \bar{\mathcal{X}}(u) + \sum_{k=1}^{\infty} \hat{\beta}_{t,k} \hat{\phi}_k(u), \quad (3.1)$$

where  $\hat{\beta}_{t,k}$  is obtained by projecting  $[\mathcal{X}_t(u) - \bar{\mathcal{X}}(u)]$  onto the direction of the  $k$ th estimated functional principal component  $\hat{\phi}_k(u)$ , that is,  $\hat{\beta}_{t,k} = \langle \mathcal{X}_t(u) - \bar{\mathcal{X}}(u), \hat{\phi}_k(u) \rangle$ .

To achieve dimensionality reduction, we truncate (3.1) to the first  $K$  terms. Because eigenvalues are ordered in a nonincreasing pattern, the first few terms often capture main mode of information in the data. Without losing much information,  $\mathcal{X}_t(u)$  can be adequately captured by the  $K$ -dimensional vector  $(\hat{\beta}_1, \hat{\beta}_2, \dots, \hat{\beta}_K)$ , where  $\hat{\beta}_1 = (\hat{\beta}_{1,1}, \hat{\beta}_{1,2}, \dots, \hat{\beta}_{1,n})$ . We obtain

$$\mathcal{X}_t(u) = \bar{\mathcal{X}}(u) + \sum_{k=1}^K \hat{\beta}_{t,k} \hat{\phi}_k(u) + e_t(u), \quad (3.2)$$

where  $e_t(u)$  represents the error terms containing estimated functional principal components and their associated scores excluded from the first  $K$  terms. The number of retained components,  $K$ , was chosen as

where  $\nu = 1/\ln[\max(\hat{\lambda}_1, n)]$  is a prespecified positive number,  $k_{\max} = \{k | \hat{\lambda}_k \geq \sum_{k=1}^n \hat{\lambda}_k/n\}$ , and  $\mathbb{1}\{\cdot\}$  represents the binary indicator function.

The sets of estimated principal component scores,  $(\hat{\beta}_1, \hat{\beta}_2, \dots, \hat{\beta}_K)$ , can be modeled and forecasted via a multivariate time-series forecasting method, such as vector autoregressive of order  $p$ , denoted by VAR( $p$ ),

$$\hat{\beta}_t = \sum_{\xi=1}^p \hat{A}_{\xi,p} \hat{\beta}_{t-\xi} + \hat{\epsilon}_t, \quad (3.4)$$

where  $t = p+1, p+2, \dots, n$ ,  $\hat{A}_{\xi,p}$  is an estimated  $(K \times K)$  coefficient matrix of the forward score series, and  $(\hat{\epsilon}_{p+1}, \hat{\epsilon}_{p+2}, \dots, \hat{\epsilon}_n)$  are the residuals obtained by fitting the VAR( $p$ ) model to the  $K$ -dimensional multivariate time series of estimated scores. The order  $p$  of a fitted VAR model can be chosen using a corrected Akaike information criterion (AIC<sub>c</sub>) of Hurvich and Tsai (1993) by minimizing

$$\text{AIC}_c(p) = n \ln |\hat{\Sigma}_{\hat{\epsilon}_p}| + \frac{n(nK + pK^2)}{n - K(p+1) - 1}, \quad (3.5)$$

over a set of possible values  $p = \{1, 2, \dots, 10\}$ , and  $\hat{\Sigma}_{\hat{\epsilon}_p} = \frac{1}{n-p} \sum_{t=p+1}^n \hat{\epsilon}_t \hat{\epsilon}_t^\top$  and  $^\top$  represents matrix transpose.

Conditional on observed time series of functions  $\mathcal{X}(u) = [\mathcal{X}_1(u), \mathcal{X}_2(u), \dots, \mathcal{X}_n(u)]$ , estimated mean function  $\bar{\mathcal{X}}(u)$  and estimated functional principal components  $\Phi(u) = [\hat{\phi}_1(u), \hat{\phi}_2(u), \dots, \hat{\phi}_K(u)]$ , the one-step-ahead point forecast can be obtained as

$$\begin{aligned} \hat{\mathcal{X}}_{n+1|n}(u) &= \mathbb{E}[\mathcal{X}_{n+1}(u)|\mathcal{X}(u), \bar{\mathcal{X}}(u), \boldsymbol{\Phi}(u)] \\ &= \bar{\mathcal{X}}(u) + \sum_{k=1}^K \hat{\beta}_{n+1|n,k} \hat{\phi}_k(u), \end{aligned}$$

where  $\hat{\phi}_k(u)$  denotes the  $k$ th estimated functional principal component and  $\hat{\beta}_{n+1|n,k}$  represents the one-step-ahead point forecasts of  $\beta_{n+1,k}$  using the VAR( $p$ ) model. When  $K = 1$ , we only have a univariate time series of principal component scores, which can be modeled and forecasted via an autoregressive (AR) model of order  $p$ .

### 3.2 | Sieve bootstrap

For every  $t \in \mathbb{Z}$ , the zero-mean random element  $\mathcal{X}_t$  is generated as

$$\mathcal{X}_t = f(\mathcal{X}_{t-1}, \mathcal{X}_{t-2}, \dots) + \varepsilon_t,$$

where  $f: \mathcal{H}^\infty \rightarrow \mathcal{H}$  is some appropriate operator and  $\{\varepsilon_t\}$  is a zero-mean independent and identically distributed (i.i.d.) innovation process in  $\mathcal{H}$  with  $\mathbb{E}\|\varepsilon_t\|^2 < \infty$  and  $\|\cdot\|$  the norm of  $\mathcal{H}$ .

Based on the last  $\ell$  observed functions,  $\mathcal{X}_{n,\ell} = (\mathcal{X}_n, \mathcal{X}_{n-1}, \dots, \mathcal{X}_{n-\ell+1})$ , for  $\ell < n$ , a predictor

$$\hat{\mathcal{X}}_{n+1} = \hat{g}(\mathcal{X}_n, \mathcal{X}_{n-1}, \dots, \mathcal{X}_{n-\ell+1}),$$

of  $\mathcal{X}_{n+1}$  is used, where  $\hat{g}: \mathcal{H}^\ell \rightarrow \mathcal{H}$  some estimated operator.

The prediction error  $\mathcal{E}_{n+1} = \mathcal{X}_{n+1} - \hat{\mathcal{X}}_{n+1}$  given  $\mathcal{X}_{n,\ell}$  can be expressed as

$$\begin{aligned} \mathcal{E}_{n+1} &= \mathcal{X}_{n+1} - \hat{\mathcal{X}}_{n+1} \\ &= \vartheta_{n+1} + [f(\mathcal{X}_n, \mathcal{X}_{n-1}, \dots) - g(\mathcal{X}_n, \mathcal{X}_{n-1}, \dots, \mathcal{X}_{n+1-\ell})] \\ &\quad + [g(\mathcal{X}_n, \mathcal{X}_{n-1}, \dots, \mathcal{X}_{n+1-\ell}) - \hat{g}(\hat{\mathcal{X}}_n, \hat{\mathcal{X}}_{n-1}, \dots, \hat{\mathcal{X}}_{n+1-\ell})] \\ &= \mathcal{E}_{I,n+1} + \mathcal{E}_{M,n+1} + \mathcal{E}_{E,n+1}, \end{aligned}$$

where  $\mathcal{E}_{I,n+1}$  is the error attributable to the i.i.d. innovation,  $\mathcal{E}_{M,n+1}$  is the model misspecification error, and  $\mathcal{E}_{E,n+1}$  is the error attributable to estimation of the unknown operator  $g$  and of the random elements  $(\mathcal{X}_n, \dots, \mathcal{X}_{n+1-\ell})$  used for one-step-ahead prediction. Both  $\mathcal{E}_{I,n+1}$  and  $\mathcal{E}_{E,n+1}$  go to zero asymptotically, but the model misspecification error  $\mathcal{E}_{M,n+1}$  may not vanish if the model used for prediction differs from the one generating the data.

We implement a sieve bootstrap method of Paparoditis and Shang (2023) to take into account all

three aforementioned sources affecting the prediction error and construct a prediction band for  $\mathcal{X}_{n+1}$  associated with the predictor  $\hat{\mathcal{X}}_{n+1}$  in a model-free way. Given  $\mathcal{X}_{n,\ell}$  of the functional time series observed and for any  $\alpha \in (0,1)$ , we construct a prediction band denoted by  $\{[\hat{\mathcal{X}}_{n+1}(u) - L_n(u), \hat{\mathcal{X}}_{n+1}(u) + U_n(u)]\}$  such that

$$\begin{aligned} \lim_{n \rightarrow \infty} \Pr(\hat{\mathcal{X}}_{n+1}(u) - L_n(u) \leq \mathcal{X}_{n+1}(u) \\ \leq \hat{\mathcal{X}}_{n+1}(u) + U_n(u), \forall u \in \mathcal{I} | \mathcal{X}_{n,\ell}) = 1 - \alpha, \end{aligned}$$

where  $\alpha$  is a level of significance.

The sieve bootstrap uses the VAR( $p$ ) process in (3.4) to generate forward score forecasts

$$\beta_{n+1}^* = \sum_{\xi=1}^p \hat{\mathbf{A}}_{\xi,p} \beta_{n+1-\xi}^* + \epsilon_{n+1}^*, \quad (3.6)$$

where  $\beta_{n+1-\xi}^* = \hat{\beta}_{n+1-\xi}$  for  $n+1-\xi \leq n$  and  $\epsilon_{n+1}^*$  is i.i.d. resampled from the set of centered residuals  $\{\hat{\epsilon}_t - \bar{\epsilon}, t = p+1, p+2, \dots, n\}$ , and  $\bar{\epsilon} = \frac{1}{n-p} \sum_{t=p+1}^n \hat{\epsilon}_t$ . From (3.2), we compute

$$\mathcal{X}_{n+1}^*(u) = \bar{\mathcal{X}}(u) + \sum_{k=1}^K \beta_{n+1,k}^* \hat{\phi}_k(u) + e_{n+1}^*(u),$$

where  $e_{n+1}^*(u)$  are i.i.d. resampled from the set  $\{e_t(u) - \bar{e}(u), t = 1, 2, \dots, n\}$  and  $\bar{e}(u) = \frac{1}{n} \sum_{t=1}^n e_t(u)$  and  $e_t(u) = \mathcal{X}_t(u) - \bar{\mathcal{X}}(u) - \sum_{k=1}^K \hat{\beta}_{t,k} \hat{\phi}_k(u)$ .

Because of stationarity, the VAR( $p$ ) process can also go backward in time to generate bootstrap samples of the estimated scores; that is,

$$\hat{\beta}_t = \sum_{\xi=1}^p \hat{\mathbf{B}}_{\xi,p} \hat{\beta}_{t+\xi} + \eta_t, \quad t = 1, 2, \dots, n-p,$$

where  $\hat{\mathbf{B}}_{\xi,p}$  denotes an estimated  $(K \times K)$  coefficient matrix for the backward score series and  $\eta_t$  is the error term, related to  $\epsilon_t$  in (3.4). The bootstrap samples  $\eta_t^*$  can be obtained by

$$\eta_t^* = \mathbf{B}_p(L^{-1}) \mathbf{A}_p^{-1}(L) \epsilon_t^*,$$

where  $\mathbf{A}_p(z) = \mathbf{I}_K - \sum_{\xi=1}^p \mathbf{A}_{\xi,p} z^\xi$ ,  $\mathbf{B}_p(z) = \mathbf{I}_K - \sum_{\xi=1}^p \mathbf{B}_{\xi,p} z^\xi$ ,  $\mathbf{I}_K$  denotes a  $(K \times K)$  diagonal matrix,  $z \in \mathbb{C}$ ,  $\mathbb{C}$  denotes a complex number, and  $\epsilon_t^*$  denotes the bootstrapped forward scores in (3.6).

The bootstrap samples for the backward scores can be obtained via



$$\beta_t^* = \sum_{\xi=1}^p \hat{B}_{\xi,p} \beta_{t+\xi}^* + \eta_t^*.$$

From (3.2), a pseudo-functional time series  $\{\mathcal{X}_1^*(u), \mathcal{X}_2^*(u), \dots, \mathcal{X}_n^*(u)\}$  can be obtained via

$$\mathcal{X}_t^*(u) = \bar{\mathcal{X}}(u) + \sum_{k=1}^K \beta_{t,k}^* \hat{\phi}_k(u) + e_t^*(u), \quad (3.7)$$

where  $e_t^*(u)$  are i.i.d. pseudo-element resampled from  $\{e_t(u) - \bar{e}(u), t = 1, 2, \dots, n\}$ .

For each pseudo-functional time series, one can apply any functional time-series forecasting method, such as the FAR(1) model of Bosq (2000), to obtain the prediction. The FAR(1) model is given by

$$\hat{\mathcal{X}}_{n+1}(u) = \bar{\mathcal{X}}(u) + \gamma[\mathcal{X}_n(u) - \bar{\mathcal{X}}(u)],$$

where  $\gamma[\cdot]$  is a bounded linear operator capturing the first-lag autocorrelation. The operator  $\gamma$  can be estimated from empirical variance and autocovariance operators

$$\begin{aligned} \hat{\gamma} &= \frac{\hat{\Gamma}(1)}{\hat{\Gamma}(0)}, \\ \hat{\Gamma}(0) &= \frac{1}{n} \sum_{t=1}^n [\mathcal{X}_t(u) - \bar{\mathcal{X}}(u)] \otimes [\mathcal{X}_t(u) - \bar{\mathcal{X}}(u)], \\ \hat{\Gamma}(1) &= \frac{1}{n} \sum_{t=1}^{n-1} [\mathcal{X}_t(u) - \bar{\mathcal{X}}(u)] \otimes [\mathcal{X}_{t+1}(u) - \bar{\mathcal{X}}(u)]. \end{aligned}$$

The distribution of prediction error  $\mathcal{E}_{n+1}^*(u) = \mathcal{X}_{n+1}^*(u) - \hat{\mathcal{X}}_{n+1}^*(u)$  is a good proxy for the distribution of  $\mathcal{E}_{n+1}(u) = \mathcal{X}_{n+1}(u) - \hat{\mathcal{X}}_{n+1}(u)$  given  $[\mathcal{X}_{n-\ell+1}(u), \dots, \mathcal{X}_n(u)]$ , where  $\hat{\mathcal{X}}_{n+1}(u)$  is the one-step-ahead point forecast obtained by the same functional time-series forecasting method, such as the FAR(1) model, applied to the original functional time series.

From  $\mathcal{E}_{n+1}^*(u)$ , we compute its standard deviation, denoted by  $\sigma_{n+1}^*(u)$ . For  $u \in \mathcal{I}$ , we compute a normalized statistic

$$V_{n+1}^*(u) = \frac{\mathcal{X}_{n+1}^*(u) - \hat{\mathcal{X}}_{n+1}^*(u)}{\sigma_{n+1}^*(u)}.$$

As shown in Paparoditis and Shang (2023), it holds true

$$\sup_{u \in \mathcal{I}} \left| [\sigma_{n+1}^*(u)]^2 - \sigma_{n+1}^2(u) \right| \rightarrow 0, \text{ in probability.}$$

Thus,  $V_{n+1}^*(u)$  is a good proxy for the distribution of  $V_{n+1}(u) = [\mathcal{X}_{n+1}(u) - \hat{\mathcal{X}}_{n+1}(u)] / \sigma_{n+1}(u)$ .

The  $(1 - \alpha)$  pointwise prediction interval for  $\mathcal{X}_{n+1}(u_i)$  can be constructed as  $\alpha/2$  and  $(1 - \alpha/2)$  quantiles of  $[\hat{\mathcal{X}}_{n+1}(u_i) + \mathcal{E}_{n+1}^*(u_i)]$  for  $i = 2, \dots, \tau$ . Let  $M^* = \sup_{u \in \mathcal{I}} |V_{n+1}^*(u)|$ , and denote by  $Q_{1-\alpha}^*$  the  $(1 - \alpha)$  quantile of the distribution of  $M^*$ . The  $(1 - \alpha)$  uniform prediction band for  $\mathcal{X}_{n+1}(u)$  associated with the predictor  $\hat{\mathcal{X}}_{n+1}(u)$  is then given by

$$[\hat{\mathcal{X}}_{n+1}(u) - Q_{1-\alpha}^* \sigma_{n+1}^*(u), \hat{\mathcal{X}}_{n+1}(u) + Q_{1-\alpha}^* \sigma_{n+1}^*(u)].$$

## 4 | UPDATING POINT AND INTERVAL FORECASTS

When a functional time series is formed as segments of a univariate time series, the most recent curve is observed sequentially and may not represent a complete curve. Denote the first  $m$  time periods of  $\mathcal{X}_{n+1}(u)$  by  $\mathcal{X}_{n+1}(u_e) = [\mathcal{X}_{n+1}(u_2), \mathcal{X}_{n+1}(u_3), \dots, \mathcal{X}_{n+1}(u_m)]^\top$ , we are particularly interested in forecasting the data in the remainder of day  $n + 1$ , denoted by  $\mathcal{X}_{n+1}(u_l)$ , where  $u_l \in (u_m, u_\tau]$ . In Figure 2, we present a conceptual diagram for dynamic updating.

The forecasting method in Section 3.1 does not use the most recent data. To improve forecast accuracy, it is desirable to dynamically update the point forecasts for the remaining period of the most recent curve  $n + 1$  by incorporating the newly arrived information (see, e.g., Chiou, 2012; Leroux et al., 2018; Yang & Yao, 2022). We consider two dynamic updating methods in Shang (2017) and Shang et al. (2019) that work well with the sieve bootstrap, which is one of the contributions of the present paper. While Shang (2017) considered dynamic updating in the S&P 500 return, Shang et al. (2019) studied it in a VIX implied by S&P 500 index option prices.




FIGURE 2 Conceptual diagram of dynamic updating.

## 4.1 | PLS

Let  $\mathcal{X}_{n+1}^c(u_e) = \mathcal{X}_{n+1}(u_e) - \bar{\mathcal{X}}(u_e)$ . This dynamic updating method shrinks regression coefficient estimates towards  $\hat{\beta}_{n+1}^{\text{TS}}$  in (3.4). The PLS regression coefficient estimates minimize a penalized residual sum of squares:

$$\arg \min_{\beta_{n+1}} \left\{ [\mathcal{X}_{n+1}^c(u_e) - \mathcal{F}_e \beta_{n+1}]^\top [\mathcal{X}_{n+1}^c(u_e) - \mathcal{F}_e \beta_{n+1}] + \lambda (\beta_{n+1} - \hat{\beta}_{n+1}^{\text{TS}})^\top (\beta_{n+1} - \hat{\beta}_{n+1}^{\text{TS}}) \right\} \quad (4.1)$$

where  $\lambda \in (0, \infty)$  is a shrinkage parameter and  $\mathcal{F}_e$  is the  $m \times K$  matrix, whose  $(i, k)$  th entry is  $\hat{\phi}_k(u_i)$  for  $2 \leq i \leq m$  and  $1 \leq k \leq K$ . The first term in (4.1) measures the in-sample goodness-of-fit, while the second term penalizes the departure of the regression coefficient estimates from  $\hat{\beta}_{n+1}^{\text{TS}}$ . By taking the first derivative with respect to  $\beta_{n+1}$ , we obtain

$$\hat{\beta}_{n+1}^{\text{PLS}} = (\mathcal{F}_e^\top \mathcal{F}_e + \lambda \mathbf{I}_K)^{-1} [\mathcal{F}_e^\top \mathcal{X}_{n+1}^c(u_e) + \lambda \hat{\beta}_{n+1}^{\text{TS}}]. \quad (4.2)$$

When the shrinkage parameter  $\lambda \rightarrow 0$ ,  $\hat{\beta}_{n+1}^{\text{PLS}}$  approach to  $\hat{\beta}_{n+1}^{\text{OLS}} = (\mathcal{F}_e^\top \mathcal{F}_e)^{-1} \mathcal{F}_e^\top \mathcal{X}_{n+1}^c(u_e)$  given  $(\mathcal{F}_e^\top \mathcal{F}_e)^{-1}$  exists; when  $\lambda \rightarrow \infty$ ,  $\hat{\beta}_{n+1}^{\text{PLS}}$  approach to  $\hat{\beta}_{n+1}^{\text{TS}}$ ; when  $0 < \lambda < \infty$ ,  $\hat{\beta}_{n+1}^{\text{PLS}}$  is a weighted average between  $\hat{\beta}_{n+1}^{\text{TS}}$  and  $\hat{\beta}_{n+1}^{\text{OLS}}$ . With an optimal selection of  $\lambda$ , the PLS forecasts of  $\mathcal{X}_{n+1}(u_i)$  at discretized time points can be given by

$$\hat{\mathcal{X}}_{n+1}^{\text{PLS}}(u_i) = \bar{\mathcal{X}}(u_i) + \sum_{k=1}^K \hat{\beta}_{n+1,k}^{\text{PLS}} \hat{\phi}_k(u_i).$$

The finite-sample performance of the PLS method crucially depends on the optimal selection of  $\lambda$ . The optimal values of  $\lambda$  need to vary intraday to achieve minimal forecast errors. To select  $\lambda$ , we consider a holdout forecast evaluation. We divide a time series of functions into a training sample consisting of curves one to 200 and a testing sample consisting of curves 201 to 250. The training sample size covers 80% of the total amount of the data, while the remaining 20% is the testing sample. The training data cover about ten months of trading, from January 4, 2021, to October 14, 2021. Within the training sample, we further divide the data into a training set consisting of Curves 1 to 150 and a validation set consisting of Curves 151 to 200. Within the validation set, the optimal values of  $\lambda$  for different updating periods were determined by

minimizing the averaged forecast error criteria in Section 5. We implicitly assume that the  $\lambda$  values selected from the validation set do not differ much from the  $\lambda$  values in the testing samples.

When we have partially observed data, we can also dynamically update prediction intervals through the sieve bootstrap in Section 3.2. We bootstrap  $B$  samples of the TS forecast regression coefficient, denoted by  $\hat{\beta}_{n+1|n}^{*,\text{TS}} = (\hat{\beta}_{n+1|n,1}^{*,\text{TS}}, \hat{\beta}_{n+1|n,2}^{*,\text{TS}}, \dots, \hat{\beta}_{n+1|n,K}^{*,\text{TS}})^\top$ . From (4.2), these bootstrap TS regression coefficient estimates lead to  $\hat{\beta}_{n+1}^{*,\text{PLS}}$ . From  $\hat{\beta}_{n+1}^{*,\text{PLS}}$ , we obtain  $B$  replications of

$$\hat{\mathcal{X}}_{n+1}^{*,\text{PLS}}(u_i) = \bar{\mathcal{X}}(u_i) + \sum_{k=1}^K \beta_{n+1,k}^{*,\text{PLS}} \hat{\phi}_k(u_i) + e_{n+1}^*(u_i),$$

where  $e_{n+1}^*(u_i)$  denotes the bootstrapped residuals corresponding to the updating period, see (3.7). The  $(1 - \alpha)$  prediction intervals for the updated forecasts are defined as  $\alpha/2$  and  $(1 - \alpha/2)$  empirical quantiles of  $\{\hat{\mathcal{X}}_{n+1}^{1,\text{PLS}}(u_i), \hat{\mathcal{X}}_{n+1}^{2,\text{PLS}}(u_i), \dots, \hat{\mathcal{X}}_{n+1}^{B,\text{PLS}}(u_i)\}$ , where  $B = 400$  denotes the number of bootstrap samples.

## 4.2 | Function-on-function linear regression

By discretizing a function into a set of grids, the PLS method updates forecasts for the updating period. Dynamic updating can also be achieved via a function-on-function linear regression of Ramsay and Silverman (2005, Chapter 16), given by

$$\mathcal{X}_{n+1}^l(u) = \bar{\mathcal{X}}^l(u) + \int_v [\mathcal{X}_{n+1}^e(v) - \bar{\mathcal{X}}^e(v)] \beta(u, v) dv + \zeta_{n+1}^l(u),$$

where  $v \in [u_2, u_m]$  and  $u \in (u_m, u_\tau]$  represent two function support ranges for the partially observed and updating data periods;  $\bar{\mathcal{X}}^e(v)$  and  $\bar{\mathcal{X}}^l(u)$  represent two mean functions for the partially observed and updating data periods;  $\mathcal{X}_{n+1}^e(v)$  and  $\mathcal{X}_{n+1}^l(u)$  denote the functional predictor and response, respectively; and  $\beta(u, v)$  denotes a bivariate regression coefficient function.

The time series of functions is divided into two parts: one block contains all data up to the most recent intraday period, and the other contains all data in the updating period. Via the functional principal component analysis, we obtain

$$\begin{aligned} \mathcal{X}_t^e(v) &= \bar{\mathcal{X}}^e(v) + \sum_{r=1}^{\infty} \hat{\theta}_{t,r} \hat{\phi}_r^e(v) \\ &= \bar{\mathcal{X}}^e(v) + \sum_{r=1}^R \hat{\theta}_{t,r} \hat{\phi}_r^e(v) + \kappa_t^e(v), \\ \mathcal{X}_t^l(u) &= \bar{\mathcal{X}}^l(u) + \sum_{s=1}^{\infty} \hat{\vartheta}_{t,s} \hat{\phi}_s^l(u) \\ &= \bar{\mathcal{X}}^l(u) + \sum_{s=1}^S \hat{\vartheta}_{t,s} \hat{\phi}_s^l(u) + \delta_t^l(u), \end{aligned}$$

where  $\hat{\phi}_r^e(v)$  and  $\hat{\phi}_s^l(u)$  denote the  $r$ th and  $s$ th estimated functional principal components associated with the partially observed and updating data periods;  $\hat{\theta}_{t,r}$  and  $\hat{\vartheta}_{t,s}$  are the principal component scores of  $\mathcal{X}_t^e(v)$  and  $\mathcal{X}_t^l(u)$ ;  $R$  and  $S$  are the retained numbers of components selected by (3.3);  $\kappa_t^e(v)$  and  $\delta_t^l(u)$  represent the error functions associated with the partially observed and updating periods.

To estimate  $\beta(u, v)$ , we link two sets of principal component scores. Let  $\hat{\theta} = [\hat{\theta}_1, \hat{\theta}_2, \dots, \hat{\theta}_R]$ , where  $\hat{\theta}_r = [\hat{\theta}_{1,r}, \hat{\theta}_{2,r}, \dots, \hat{\theta}_{n,r}]^\top$  and  $\hat{\vartheta} = [\hat{\vartheta}_1, \hat{\vartheta}_2, \dots, \hat{\vartheta}_S]$ , where  $\hat{\vartheta}_s = [\hat{\vartheta}_{1,s}, \hat{\vartheta}_{2,s}, \dots, \hat{\vartheta}_{n,s}]^\top$ . Via ordinary least squares, we study a linear relationship between  $\hat{\theta}$  and  $\hat{\vartheta}$ , given by

$$\hat{\vartheta} = \hat{\theta} \times \rho, \tag{4.3}$$

where  $\rho$  can be estimated by

$$\hat{\rho} = (\hat{\theta}^\top \hat{\theta})^{-1} \hat{\theta}^\top \hat{\vartheta}. \tag{4.4}$$

The one-step-ahead point forecast of  $\mathcal{X}_{n+1}^l(u)$  can be obtained by

$$\hat{\mathcal{X}}_{n+1}^l(u) = \bar{\mathcal{X}}^l(u) + \sum_{s=1}^S \hat{\vartheta}_{n+1,s} \hat{\phi}_s^l(u). \tag{4.5}$$

Given (4.3), (4.5) can be approximated by

$$\hat{\mathcal{X}}_{n+1}^l(u) \approx \bar{\mathcal{X}}^l(u) + \hat{\theta}_{n+1} \times \hat{\rho} \times \hat{\phi}^l(u), \tag{4.6}$$

where  $\hat{\phi}^l(u) = [\hat{\phi}_1^l(u), \hat{\phi}_2^l(u), \dots, \hat{\phi}_S^l(u)]$ .

The one-step-ahead interval forecast of  $\mathcal{X}_{n+1}^l(u)$  can be obtained by

$$\begin{aligned} \hat{\mathcal{X}}_{n+1}^{l,*}(u) &= \bar{\mathcal{X}}^l(u) + \int_{v \in \mathcal{I}_e} [\mathcal{X}_{n+1}^e(v) - \bar{\mathcal{X}}^e(v)] \hat{\beta}^*(u, v) dv \\ &\quad + e_{n+1}^{l,*}(u), \end{aligned}$$

where  $\hat{\beta}^*(u, v)$  represents the bootstrapped regression coefficient estimates and  $e_{n+1}^{l,*}(u)$  represents the bootstrapped residuals corresponding to the updating period, see (3.7). Via the sieve bootstrap, we obtain bootstrap curves. By projecting these curves onto the estimated functional principal components, we obtain bootstrap principal component scores. Via (4.4), we obtain bootstrap  $\rho^*$ , from which we form the bootstrap forecasts for the updating period via (4.6).

The  $(1 - \alpha)$  prediction intervals for the updated forecasts are defined as  $\alpha/2$  and  $(1 - \alpha/2)$  empirical quantiles of  $\{\hat{\mathcal{X}}_{n+1}^{l,1}(u), \hat{\mathcal{X}}_{n+1}^{l,2}(u), \dots, \hat{\mathcal{X}}_{n+1}^{l,B}(u)\}$ .

## 5 | EVALUATIONS OF POINT AND INTERVAL FORECAST ACCURACY

In the forecasting literature, two schemes exist to evaluate forecast accuracy: the expanding and rolling windows. While a rolling window is a fixed size, expanding window has a fixed starting point and incorporates the most recent data as it becomes available. We consider the expanding-window scheme. The initial training samples are curves from Days 1 to 200, and we compute a 1-day-ahead forecast using the time-series method for day 201. Through an expanding-window approach, we increase the training samples from Days 1 to 201, and we then compute the 1-day-ahead forecast for Day 202. We iterate this procedure until the training samples cover the entire 250 days.

### 5.1 | Mean squared forecast error (MSFE)

We compute the out-of-sample point forecasts and evaluate their accuracy by MSFE. The MSFE measures the closeness of the forecasts compared with the actual values of the variable being forecast.

$$\begin{aligned} \text{MSFE}(u_i) &= \frac{1}{n_{\text{test}}} \sum_{i=1}^{n_{\text{test}}} [\mathcal{X}_i(u_i) - \hat{\mathcal{X}}_i(u_i)]^2 \\ \text{MSFE} &= \frac{1}{\tau - 1} \sum_{i=2}^{\tau} \text{MSFE}(u_i) \end{aligned}$$

where  $\mathcal{X}_i(u_i)$  represents the actual holdout samples for the  $i$ th intraday period on the  $i$ th day,  $\hat{\mathcal{X}}_i(u_i)$  represents the one-step-ahead point forecasts for the holdout samples, and  $n_{\text{test}} = 50$  represents the number of curves in the out-of-sample forecasting period.



## 5.2 | Empirical coverage probability and interval score

To evaluate interval forecast accuracy, we compute empirical coverage probabilities (ECPs) for pointwise prediction intervals and uniform prediction bands. The ECPs are defined by

$$\begin{aligned} \text{ECP}_{\text{pointwise}} &= 1 - \frac{1}{n_{\text{test}} \times (\tau - 1)} \sum_{i=1}^{n_{\text{test}}} \sum_{i=2}^{\tau} \left[ \mathbb{1} \left\{ \mathcal{X}_i(u_i) < \hat{\mathcal{X}}_i^{\text{lb}}(u_i) \right\} + \mathbb{1} \left\{ \mathcal{X}_i(u_i) > \hat{\mathcal{X}}_i^{\text{ub}}(u_i) \right\} \right] \\ \text{ECP}_{\text{uniform}} &= 1 - \frac{1}{n_{\text{test}}} \sum_{i=1}^{n_{\text{test}}} \left[ \mathbb{1} \left\{ \mathcal{X}_i(u) < \hat{\mathcal{X}}_i^{\text{lb}}(u) \right\} + \mathbb{1} \left\{ \mathcal{X}_i(u) > \hat{\mathcal{X}}_i^{\text{ub}}(u) \right\} \right]. \end{aligned}$$

The uniform prediction bands are wider than the pointwise prediction intervals. As a result, the uniform empirical coverage probability counts the number of days in the forecasting period, where the constructed prediction bands cover all future realizations.

For pointwise prediction intervals, we also utilize the interval score of Gneiting and Raftery (2007) and Gneiting and Katzfuss (2014). For each year in the forecasting period, the one-step-ahead prediction intervals were computed at the  $(1 - \alpha)$  nominal coverage probability. Let  $\hat{\mathcal{X}}_i^{\text{lb}}(u_i)$  and  $\hat{\mathcal{X}}_i^{\text{ub}}(u_i)$  be the lower and upper bounds of the pointwise prediction intervals, respectively. A scoring rule for the prediction interval at time point  $u_i$  is

$$\begin{aligned} S_{\alpha} \left[ \hat{\mathcal{X}}_i^{\text{lb}}(u_i), \hat{\mathcal{X}}_i^{\text{ub}}(u_i), \mathcal{X}_i(u_i) \right] &= \left[ \hat{\mathcal{X}}_i^{\text{ub}}(u_i) - \hat{\mathcal{X}}_i^{\text{lb}}(u_i) \right] + \frac{2}{\alpha} \left[ \hat{\mathcal{X}}_i^{\text{lb}}(u_i) - \mathcal{X}_i(u_i) \right] \mathbb{1} \left\{ \mathcal{X}_i(u_i) < \hat{\mathcal{X}}_i^{\text{lb}}(u_i) \right\} \\ &\quad + \frac{2}{\alpha} \left[ \mathcal{X}_i(u_i) - \hat{\mathcal{X}}_i^{\text{ub}}(u_i) \right] \mathbb{1} \left\{ \mathcal{X}_i(u_i) > \hat{\mathcal{X}}_i^{\text{ub}}(u_i) \right\}. \end{aligned}$$

Averaged over different days in the forecasting period, the mean interval score for each intraday period  $i$  and its averaged mean interval score are defined by

$$\begin{aligned} \bar{S}_{\alpha}(u_i) &= \frac{1}{n_{\text{test}}} \sum_{i=1}^{n_{\text{test}}} S_{\alpha} \left[ \hat{\mathcal{X}}_i^{\text{lb}}(u_i), \hat{\mathcal{X}}_i^{\text{ub}}(u_i), \mathcal{X}_i(u_i) \right], \\ \bar{S}_{\alpha} &= \frac{1}{\tau - 1} \sum_{i=2}^{\tau} \bar{S}_{\alpha}(u_i). \end{aligned}$$

## 6 | RESULTS

The functional time-series forecasting method in Section 3 decomposes a time series of functions, that is, CIDRs from January 4 to December 22, 2021, into a set of

functional principal components and their associated principal component scores. The temporal dependence of the functional time series is inherited by the correlation among the principal component scores. From (3.3), the retained number of components is one and the selected order of the VAR model is  $p = 4$  in (3.5). We display and attempt to interpret the first functional principal compo-

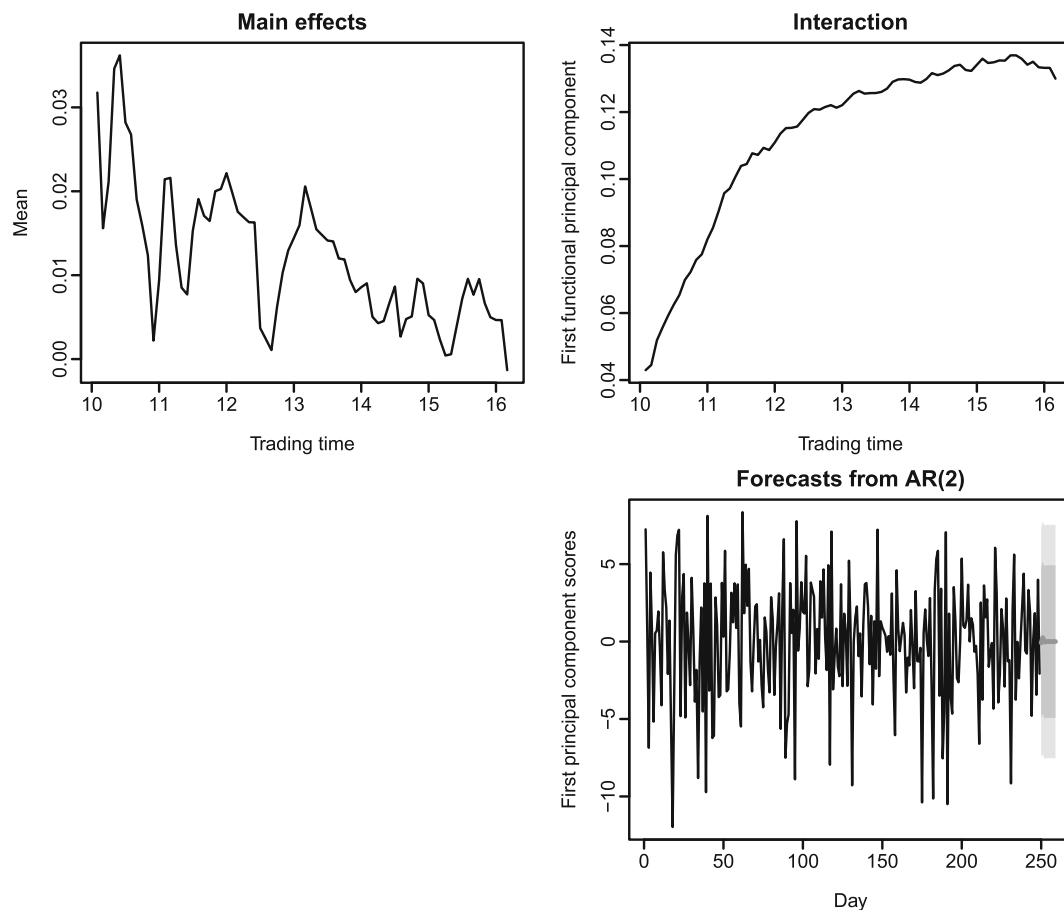
nent in the top panel of Figure 3. The mean function illustrates average changes in CIDRs with larger fluctuations at or before 14:00 and relatively smaller fluctuations after 14:00. The first functional principal component shows a general increasing trend in CIDRs. As a demonstration, the 10-day-ahead point forecasts of the principal component scores show an increasing trend for a short-time interval before stabilizing. This indicates the short-term predictive ability of the CIDRs.

To assess overall goodness-of-fit, the residual functions of the fitted functional time series are displayed in Figure 4 via a filled contour plot. As no clear pattern is exhibited in the contour plot, the fitted model captures the temporal structure well. Further, we implement a

white noise test of Gabrys and Kokoszka (2007) to check if there is any temporal correlation in the residual functions. From the  $p$ -value of 0.4453, we conclude that the residuals functions are white noise.

As an illustration, we model CIDRs from January 4 to December 22, 2021, and forecast the CIDR for December 23, 2021. We multiply the 1-day-ahead forecast principal component scores by the estimated functional principal components and add the estimated mean function. Using the sieve bootstrap method described in Section 3.2, the 80% and 95% pointwise prediction intervals and uniform prediction bands, shown in Figure 5, are constructed from these bootstrapped forecasts.

Averaged over 50 days in the forecasting period and 73 different intraday updating periods, we present the



**FIGURE 3** Estimated mean function, first set of functional principal component, and its associated principal component scores for the CIDRs of the XAO from January 4 to December 22, 2021. For the 10-day-ahead forecasts, the 80% and 95% prediction intervals of the estimated principal component scores are shown by the dark and light gray regions. For the univariate time series of the principal component scores, the AR(2) is selected as the optimal model based on the Akaike information criterion.

averaged MSFE,  $ECP_{\text{pointwise}, 1-\alpha}$ ,  $ECP_{\text{uniform}, 1-\alpha}$ , and  $\bar{S}_\alpha$  obtained from the TS method. In Table 1, the sieve bootstrap method produces the empirical coverage probabilities that are no less than the nominal coverage probabilities.

## 6.1 | Updating point forecasts


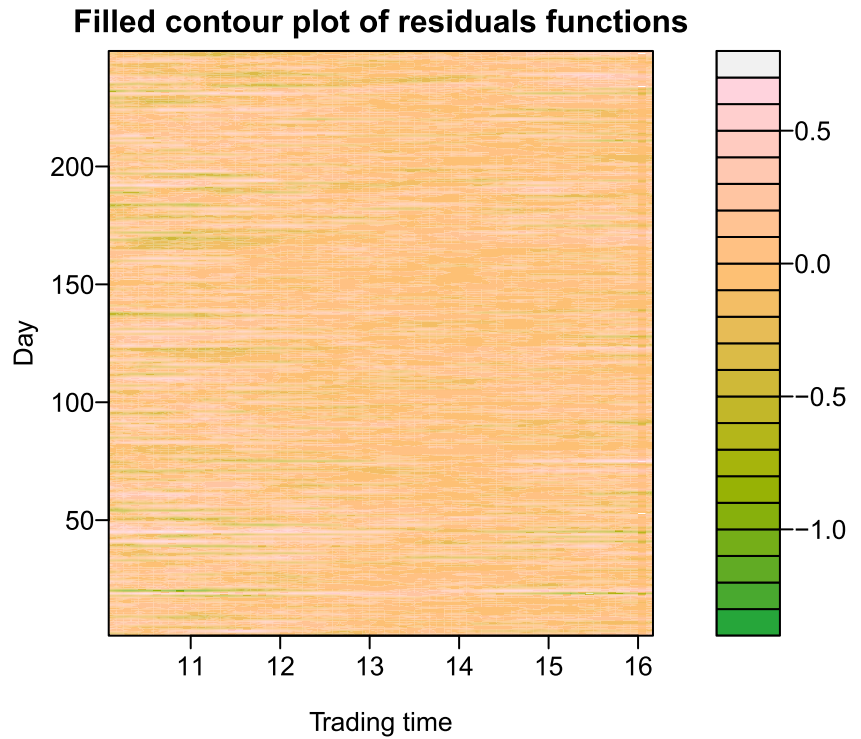
When we have partially observed data in the most recent curve, we can dynamically update our point forecasts for the remaining period to improve forecast accuracy. Two dynamic updating methods were presented in Section 4, and their point forecast accuracies are evaluated and compared in Figure 6b. The forecast errors generally decrease throughout the day for the dynamic updating methods, as we have more partially observed data in the most recent curve. Averaged over 50 days in the forecasting period, the PLS method has the best point forecast accuracy, measured by MSFE, over all evaluation intraday time points. When  $\lambda = 0$  in (4.2), the PLS method

reduces to the OLS method. The PLS forecasts produce smaller errors than the OLS forecasts in the early trading time, and then, their forecasts and forecast errors are similar in the afternoon. This phenomenon is reflected by the estimated optimal values of  $\lambda$ . As shown in Figure 6a, the  $\lambda$  values are selected by minimizing the MSFE within the validation set.

Future values of the CIDRs are hard to predict accurately, but the signs of future values are relatively easier (see also Hausman et al., 1992). Averaged over 50 days in the forecasting period, we count the number of times where forecasts and holdout values share the same sign, that is, negative or positive. Using the PLS method, we compute the probabilities of making correct sign predictions over different updating periods in Figure 7.

The probability of correct sign prediction increases from 0.7263 at the beginning of the updating period (i.e., 10:15–16:10) to 0.8200 at the end of the intraday trading time (i.e., 16:10), a peak at 0.9008 occurring at 14:10–16:10. The peak occurs when the signs of the remaining observations are most accurately predicted

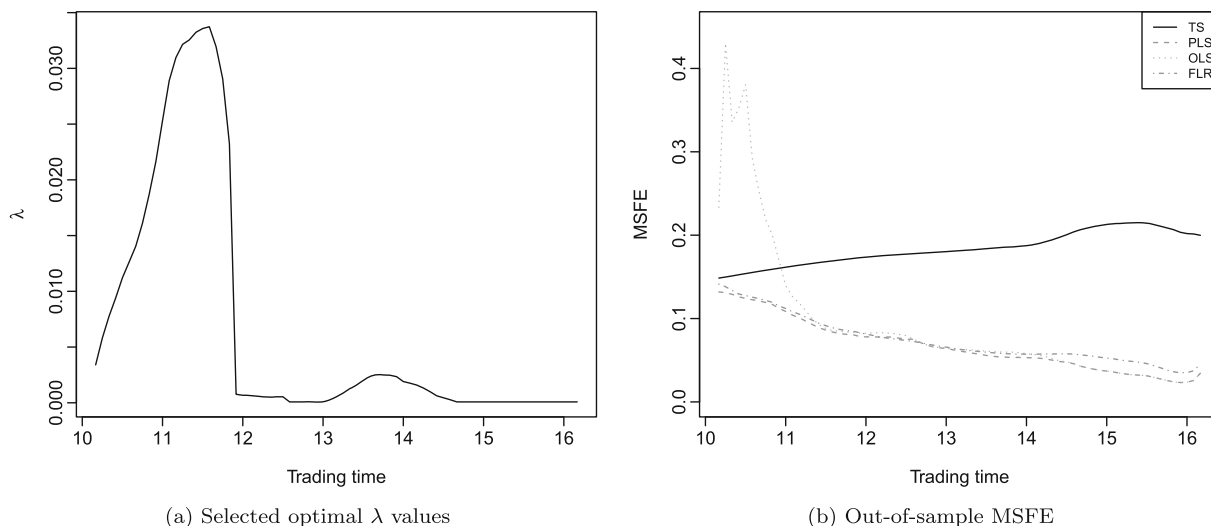
**FIGURE 4** Filled contour plot of the residual functions obtained from the fitted functional time-series model. Using a terrain color palette, negative residuals are shown in green, positive residuals in gray and pink, and residuals around zero are shown in orange.




**FIGURE 5** One-step-ahead point forecasts of the CIDRs of the XAO on December 23, 2021, the 80% and 95% pointwise prediction intervals, and uniform prediction bands constructed via the sieve bootstrap.

**TABLE 1** Averaged forecast errors of the TS method over 50 days in the forecasting period.

Method	MSFE	ECP <sub>pointwise, 1-α</sub>		ECP <sub>uniform, 1-α</sub>		$\bar{S}_α$	
		α = 0.2	α = 0.05	α = 0.2	α = 0.05	α = 0.2	α = 0.05
TS	0.1474	0.89	0.98	0.88	0.96	1.43	2.08



**FIGURE 6** A comparison of point forecast accuracy, as measured by averaged  $MSFE_i$  for  $i = 3, 4, \dots, \tau$  over 50 days in the forecasting period, among the TS, PLS, OLS, and FLR methods. We observe at least one data point in the dynamic updating methods before updating forecasts.



**FIGURE 7** Averaged over 50 days in the forecasting period, the probabilities of predicting the correct sign are computed for the PLS and FLR methods at different updating periods.

on average. We also present the correct sign prediction probability using the FLR method, which generally performs inferior to the PLS method. A plausible explanation for the inferior performance of the FLR method is that it does not use the TS regression coefficient as a priori.

As a trading strategy, a day trader can enter the market at around 14:10 because we can predict the direction more accurately. For example, if the predictions indicate the XAO will likely go up, the day trader may choose to follow the momentum and sell it at or close to the closing trading time of the day. The daily profit can be made so

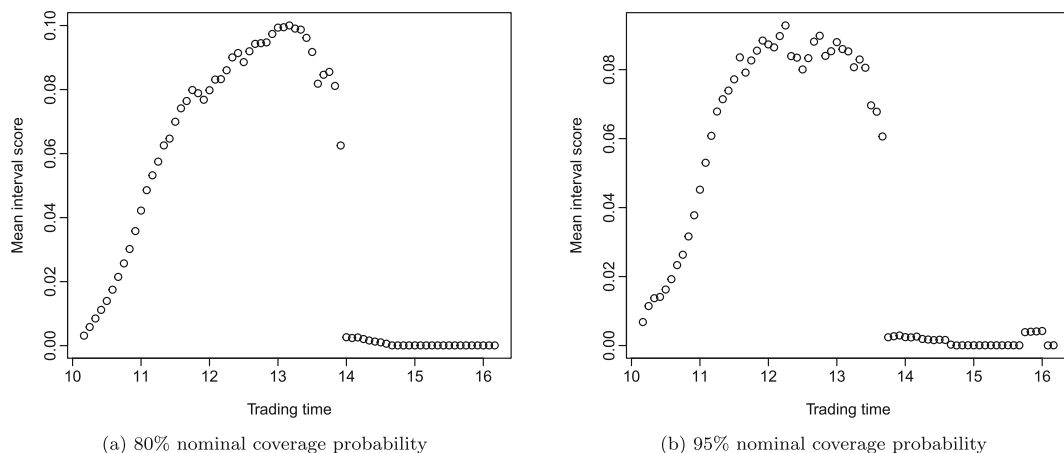
long the difference between purchasing and selling prices covers transaction costs.

## 6.2 | Updating interval forecasts

We can also update our pointwise interval forecasts for the remaining period to achieve improved accuracy when we have partially observed data in the most recent curve. We estimate the optimal values of  $\lambda$  for different updating periods by minimizing the averaged mean interval score within the validation set for the PLS method. The optimal values are presented in Figure 8 for the 80% and 95% nominal coverage probabilities. At the 80% nominal coverage probability, the values of  $\lambda$  increase from the beginning of the trading time to around 13:15, then decrease sharply from 14:00 to 16:10. At the 95% nominal coverage probability, the values of  $\lambda$  increase from the beginning of the trading time to 12:15, stabilize until 13:30, and decrease sharply from 13:40 to 16:10.

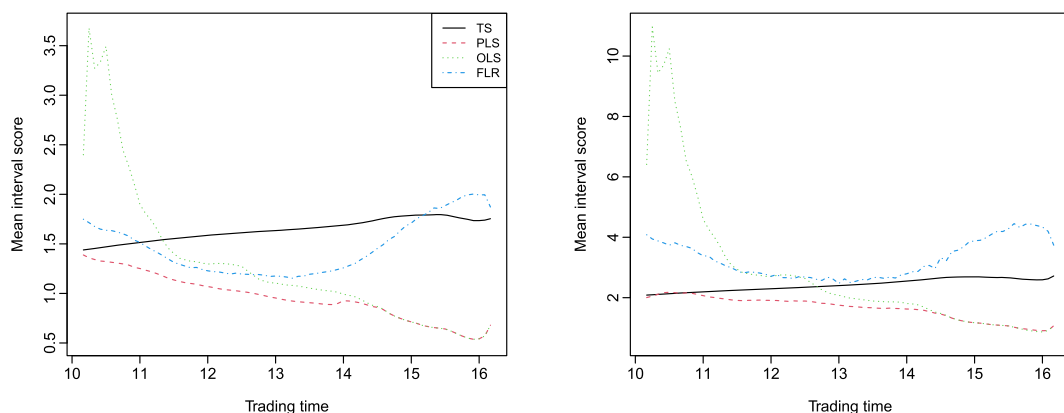
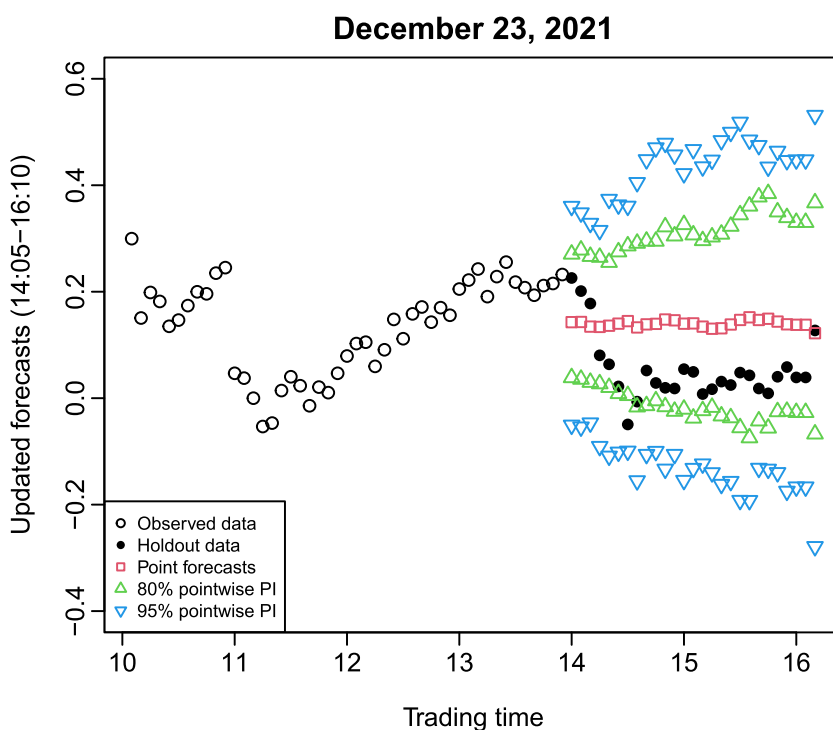
Suppose the CIDRs of the XAO have been observed from 10:05 to 14:00 on December 23, 2021, we apply the PLS and FLR methods to update the interval forecasts for the remaining period of that day. Figure 9 presents the 80% and 95% pointwise prediction intervals when the dynamic updating methods are used for the remaining period.

Between the PLS and FLR methods, as shown in Figure 10, the PLS method produces the smaller mean interval score, averaged over 50 days in the forecasting period. Thus, it provides the most accurate evaluation of forecast uncertainty.



**FIGURE 8** Estimated optimal values of  $\lambda$  for different intraday updating periods by minimizing the mean interval scores within the validation set.

**FIGURE 9** As we sequentially observe new information from the beginning of the trading time to 14:00, we apply the PLS method with optimally tuned parameters to update point and interval forecasts of the CIDRs of the XAO on December 23, 2021.



**FIGURE 10** A comparison of interval forecast accuracy, as measured by averaged mean interval score for different updating periods over 50 days in the forecasting period, among the TS, PLS, OLS, and FLR methods.



**TABLE 2** By averaging 73 different intraday updating periods, we compute the averaged forecast errors of the TS, PLS, OLS, and FLR methods over 50 days in the forecasting period.

Method	MSFE	ECP <sub>pointwise,1-<math>\alpha</math></sub>		$\bar{S}_\alpha$	
		$\alpha = 0.2$	$\alpha = 0.05$	$\alpha = 0.2$	$\alpha = 0.05$
TS	0.1838	0.8817	0.9688	1.6429	2.4277
PLS	0.0672	0.7275	0.8851	0.9564	1.6440
FLR	0.0735	0.5001	0.7266	1.4661	3.2729

Averaged over 50 days in the forecasting period and 73 different intraday updating periods, we present the averaged MSFE, ECP<sub>pointwise,1- $\alpha$</sub> , ECP<sub>uniform,1- $\alpha$</sub> , and  $\bar{S}_\alpha$  obtained from the TS, PLS, OLS, and FLR methods. In Table 2, the dynamic updating methods, such as the PLS and FLR, improve the point forecast accuracy as measured by the MSFE. In terms of the interval forecast accuracy, the ECPs of the PLS method are smaller than the nominal ones, which may be due to the same estimated values of  $\lambda$  are used in each bootstrap sample. The lack of uncertainty associated with  $\lambda$  stems from its computational cost and trade-off between accuracy and practicality. Nevertheless, the PLS method produces the smallest MSFE and  $\bar{S}_\alpha$ , and thus, it is our recommended method.

## 7 | CONCLUSION

Our forecasting method treats the observations as realizations of a functional time series, where the temporal dependence among curves can be modeled and forecasted via functional principal component analysis. The temporal dependence in the functional time series is manifested in the estimated principal component scores. By explaining at least 85% of the total variation in the data, we select the retained number of components. When the number of components is one, a multivariate time-series technique, such as VAR( $p$ ) model, is reduced to a univariate time-series technique, such as AR( $p$ ) model. Conditioning on the estimated functional principal components and historical curves, a one-step-ahead point forecast can be obtained.

When partial data in the most recent curve are sequentially observed, two dynamic updating methods, the PLS and FLR methods, are implemented to update forecasts. The FLR method decomposes two blocks of functions via functional principal component analyses. The first block is associated with the data in the partially observed period of the most recent curve; the second block is associated with the data in the updating period. By linking these estimated principal component scores

via linear regression, a regression coefficient estimate for these scores is obtained. Conditioning on the estimated functional principal components and observed data in the updating period, updated forecasts can be obtained based on the estimated regression coefficient of the scores and the principal component scores in the partially observed data of the most recent curve. The PLS method is a shrinkage estimator: When the newly arrived data occur in the morning of the trading time, it shrinks its regression coefficient towards time-series regression coefficient with a comparably larger tuning parameter  $\lambda$ . When the newly arrived data cover most of the day, it shrinks its regression coefficient towards the OLS regression coefficient. By selecting the optimal values of  $\lambda$ , the PLS method achieves superior forecast accuracy among all methods considered for producing point and interval forecasts, measured by averaged MSFE and interval score, respectively.

As a means of measuring forecast uncertainty, a sieve bootstrap method was used to construct pointwise prediction intervals and uniform prediction bands for the TS method. The sieve bootstrap method takes into account model misspecification error and thus achieves excellent calibration, where the empirical coverage probability is no less than the nominal coverage probability. The sieve bootstrap plays an important role in constructing pointwise prediction intervals for the updated forecasts using the PLS and FLR methods.

There are at least three ways in which the present paper can be extended: (1) While the methods apply to 5-min intervals, future research could consider optimal sampling frequency, possibly even shorter time intervals. (2) The presence of outliers can affect the estimation of the covariance structure. Because the covariance relies on the sample mean, it affects the estimation accuracy of the functional principal components. Consequently, it may affect the accuracy of forecasts, and one could consider robust functional principal component analysis (see, e.g., Bali et al., 2011). (3) With a set of validation samples, the PLS tuning parameters can be adaptively chosen without recomputing them.

## ACKNOWLEDGMENTS

The authors thank a reviewer for insightful comments and suggestions. This study was partially supported by an Australian Research Council Discovery Project (ARC DP230102250). The authors acknowledge seminar participants at the Department of Applied Finance, Macquarie University, for their insightful comments and suggestions. Open access publishing facilitated by Macquarie University, as part of the Wiley - Macquarie University agreement via the Council of Australian University Librarians.

## DATA AVAILABILITY STATEMENT

The data that support the findings of this study are available from the corresponding author upon request.

## ORCID

Han Lin Shang  <https://orcid.org/0000-0003-1769-6430>

## REFERENCES

- Andersen, T. G. (2000). Some reflections on analysis of high-frequency data. *Journal of Business and Economic Statistics*, 18(2), 146–153.
- Aue, A., Norinho, D. D., & Hörmann, S. (2015). On the prediction of stationary functional time series. *Journal of the American Statistical Association: Theory and Methods*, 110(509), 378–392.
- Bali, J. L., Boente, G., Tyler, D. E., & Wang, J.-L. (2011). Robust functional principal components: A projection-pursuit approach. *The Annals of Statistics*, 39(6), 2852–2882.
- Bosq, D. (2000). *Linear processes in function spaces*: Springer.
- Box, G. E. P. (1976). Science and statistics. *Journal of the American Statistical Association: Applications*, 71(356), 791–799.
- Campbell, J. Y., Lo, A. W., & MacKinlay, A. C. (1997). *The econometrics of financial markets*: Princeton University Press.
- Chiou, J.-M. (2012). Dynamical functional prediction and classification with application to traffic flow prediction. *The Annals of Applied Statistics*, 6(4), 1588–1614.
- Dacorogna, M. M., Gençay, R., Olsen, R. B., Müller, U. A., & Pictet, O. V. (2001). *An introduction to high-frequency finance*: Academic Press.
- Engle, R. F. (2000). The econometrics of ultra-high-frequency data. *Econometrica*, 68(1), 1–22.
- Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, 25(2), 383–417.
- Gabrys, R., & Kokoszka, P. (2007). Portmanteau test of independence for functional observations. *Journal of the American Statistical Association: Theory and Methods*, 102(480), 1338–1348.
- Gao, L., Han, Y., Li, S. Z., & Zhou, G. (2017). Intraday momentum: The first half-hour return predicts the last half-hour return: SSRN Working paper. [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=3D2552752](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3D2552752)
- Ghysels, E. (2000). Some econometrics recipes for high-frequency data cooking. *Journal of Business and Economic Statistics*, 18(2), 154–163.
- Gneiting, T., & Katzfuss, M. (2014). Probabilistic forecasting. *Annual Review of Statistics and Its Application*, 1, 125–151.
- Gneiting, T., & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477), 359–378.
- Goodhart, C. A. E., & O'Hara, M. (1997). High frequency data in financial markets: Issues and applications. *Journal of Empirical Finance*, 4(2–3), 73–114.
- Gouriéroux, C., & Jasiak, J. (2001). *Financial econometrics: Problems, models and methods*: Princeton University Press.
- Hansen, P. R., & Lunde, A. (2006). Realized variance and market microstructure noise. *Journal of Business & Economic Statistics*, 24(2), 127–161.
- Hausman, J. A., Lo, A. W., & MacKinlay, A. C. (1992). An ordered probit analysis of transaction stock prices. *Journal of Financial Econometrics*, 31(3), 319–379.
- Hays, S., Shen, H., & Huang, J. Z. (2012). Functional dynamic factor models with application to yield curve forecasting. *The Annals of Applied Statistics*, 6(3), 870–894.
- Heston, S. L., Korajczyk, R. A., & Sadka, R. (2010). Intraday pattern in the cross-section of stock returns. *Journal of Finance*, 65(4), 1369–1407.
- Hooker, G., & Shang, H. L. (2022). Selecting the derivative of a functional covariate in scalar-on-function regression. *Statistics and Computing*, 32(3), 35.
- Horváth, L., Kokoszka, P., & Rice, G. (2014). Testing stationarity of functional time series. *Journal of Econometrics*, 179(1), 66–82.
- Hurvich, C. M., & Tsai, C.-L. (1993). A corrected Akaike information criterion for vector autoregressive model selection. *Journal of Time Series Analysis*, 14(3), 271–279.
- Klepsch, J., & Klüppelberg, C. (2017). An innovations algorithm for the prediction of functional linear processes. *Journal of Multivariate Analysis*, 155, 252–271.
- Klepsch, J., Klüppelberg, C., & Wei, T. (2017). Prediction of functional ARMA processes with an application to traffic data. *Econometrics and Statistics*, 1, 128–149.
- Kokoszka, P., & Reimherr, M. (2013). Determining the order of the functional autoregressive model. *Journal of Time Series Analysis*, 34(1), 116–129.
- Kokoszka, P., Rice, G., & Shang, H. L. (2017). Inference for the autocovariance of a functional time series under conditional heteroscedasticity. *Journal of Multivariate Analysis*, 162, 32–50.
- Kokoszka, P., & Zhang, X. (2012). Functional prediction of intraday cumulative returns. *Statistical Modelling*, 12(4), 377–398.
- Leroux, A., Xiao, L., Crainiceanu, C., & Checkley, W. (2018). Dynamic prediction in functional concurrent regression with an application to child growth. *Statistics in Medicine*, 37(8), 1376–1388.
- Li, D., Robinson, P. M., & Shang, H. L. (2020). Long-range dependent curve time series. *Journal of the American Statistical Association: Theory and Methods*, 115(530), 957–971.
- Li, D., Robinson, P. M., & Shang, H. L. (2021). Local Whittle estimation of long-range dependence for functional time series. *Journal of Time Series Analysis*, 42(5–6), 685–695.
- Li, D., Robinson, P. M., & Shang, H. L. (2023). Nonstationary fractionally integrated functional time series. *Bernoulli*, 29(2), 1505–1526.
- Lintner, J. (1965). The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *The Review of Economics and Statistics*, 47(1), 13–37.
- Lyons, R. (2001). *The microstructure approach to exchange rates*: MIT Press.
- Papadimitis, E., & Shang, H. L. (2023). Bootstrap prediction bands for functional time series. *Journal of the American Statistical Association: Theory and Methods*, 118(542), 972–986.
- Peña, D., & Tsay, R. S. (2021). *Statistical learning for big dependent data*: John Wiley & Sons, Inc.
- Ramsay, J. O., & Silverman, B. W. (2005). *Functional data analysis* (2nd ed.). Springer.
- Rice, G., Wirjanto, T., & Zhao, Y. (2020). Tests for conditional heteroscedasticity of functional data. *Journal of Time Series Analysis*, 41(6), 733–768.

- Shang, H. L. (2017). Forecasting intraday S&P 500 index returns: A functional time series approach. *Journal of Forecasting*, 36(7), 741–755.
- Shang, H. L. (2019). Visualizing rate of change: An application to age-specific fertility rates. *Journal of the Royal Statistical Society: Series A*, 182(1), 249–262.
- Shang, H. L., Yang, Y., & Kearney, F. (2019). Intraday forecasts of a volatility index: Functional time series methods with dynamic updating. *Annals of Operations Research*, 282, 331–354.
- Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance*, 19(3), 425–442.
- Tsay, R. S. (2010). *Analysis of financial time series* (3rd ed.). John Wiley & Sons, Inc.
- Wang, S., Jank, W., & Shmueli, G. (2008). Explaining and forecasting online auction prices and their dynamics using functional data analysis. *Journal of Business and Economic Statistics*, 26(2), 144–160.
- Wood, R. A. (2000). Market microstructure research databases: History and projections. *Journal of Business and Economic Statistics*, 18(2), 140–145.
- Yang, Y., & Yao, F. (2022). Online estimation for functional data. *Journal of the American Statistical Association: Theory and Methods*.
- Zhang, X., & Wang, J-L. (2016). From sparse to dense functional data and beyond. *The Annals of Statistics*, 44(5), 2281–2321.

## AUTHOR BIOGRAPHIES

**Dr Han Lin Shang** holds a PhD in Statistics from Monash University. He is a professor of Business Analytics at Macquarie University.

**Dr Kaiying Ji** is a lecturer in the Discipline of Accounting, Governance and Regulation at the University of Sydney Business School. Kaiying is a member of CPA Australia and the CFA Institute. She has research interests in forecasting, financial accounting regulation and disclosure, capital market, and asset pricing.

**How to cite this article:** Shang, H. L., & Ji, K. (2023). Forecasting intraday financial time series with sieve bootstrapping and dynamic updating. *Journal of Forecasting*, 42(8), 1973–1988. <https://doi.org/10.1002/for.3000>