



**QUEEN'S
UNIVERSITY
BELFAST**

Intraday Forecasts of a Volatility Index: Functional Time Series Methods with Dynamic Updating

Shang, H. L., Yang, Y., & Kearney, F. (2018). Intraday Forecasts of a Volatility Index: Functional Time Series Methods with Dynamic Updating. *Annals of Operations Research*. Advance online publication. <https://doi.org/10.1007/s10479-018-3108-4>

Published in:
Annals of Operations Research

Document Version:
Peer reviewed version

Queen's University Belfast - Research Portal:
[Link to publication record in Queen's University Belfast Research Portal](#)

Publisher rights
© 2018 Springer Science+Business Media, LLC, part of Springer Nature.
This work is made available online in accordance with the publisher's policies. Please refer to any applicable terms of use of the publisher

General rights
Copyright for the publications made accessible via the Queen's University Belfast Research Portal is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy
The Research Portal is Queen's institutional repository that provides access to Queen's research output. Every effort has been made to ensure that content in the Research Portal does not infringe any person's rights, or applicable UK laws. If you discover content in the Research Portal that you believe breaches copyright or violates any law, please contact openaccess@qub.ac.uk.

Open Access
This research has been made openly available by Queen's academics and its Open Research team. We would love to hear how access to this research benefits you. – Share your feedback with us: <http://go.qub.ac.uk/oa-feedback>

Intraday forecasts of a volatility index: Functional time series methods with dynamic updating

Han Lin Shang*, Yang Yang,
Research School of Finance, Actuarial Studies and Statistics
Australian National University

Fearghal Kearney
Queen's Management School
Queen's University Belfast

October 12, 2018

Abstract

As a forward-looking measure of future equity market volatility, the VIX index has gained immense popularity in recent years to become a key measure of risk for market analysts and academics. We consider discrete reported intraday VIX tick values as realisations of a collection of curves observed sequentially on equally spaced and dense grids over time and utilise functional data analysis techniques to produce one-day-ahead forecasts of these curves. The proposed method facilitates the investigation of dynamic changes in the index over very short time intervals as showcased using the 15-second high-frequency VIX index values. With the help of dynamic updating techniques, our point and interval forecasts are shown to enjoy improved accuracy over conventional time series models.

Keywords: functional principal component regression; functional linear regression; ordinary least squares; penalise least squares; high-frequency financial data.

*Postal address: Research School of Finance, Actuarial Studies and Statistics, Level 4, Building 26C, Australian National University, Kingsley Street, Canberra, ACT 2601, Australia; Telephone: +61(2) 612 50535; Fax: +61(2) 612 50087; Email: hanlin.shang@anu.edu.au.

1 Introduction

In 1993, the Chicago Board Options Exchange (CBOE) introduced the CBOE volatility index (VIX) implied by S&P 500 Index option prices. VIX is computed based on a weighted portfolio of 30-day S&P 500 calls and puts, to construct a forward-looking measure of future equity market volatility. Driven by the introduction of a range of VIX derivatives: futures, options, and exchange-traded notes, the VIX index has gained immense popularity in recent years. It has become a key measure of risk for market analysts and academics alike, causing it to be referred to as the ‘investor fear gauge’ by the financial press and investment community (see, e.g., [Whaley 2000](#)). Therefore, producing accurate forecasts of future VIX index values are of great importance to risk management. Such forecasts are of interest to all parties involved in analysing VIX and VIX-related instruments, including hedge and pension funds, endowments, and even retail investors.

While producing volatility forecasts of the S&P 500 and other equity indices is popular in the literature (see, e.g., [Wang and Zhu 2010](#), [Xu 1999](#), [Bali and Theodossiou 2007](#), [Kılıç and Uğur 2018](#)), modelling implied volatility indices have received less research attention. Researchers have previously considered VIX index values as discrete time series, producing forecasts using autoregressive fractionally integrated moving average models and heterogeneous autoregressive models (see, e.g., [Konstantinidi et al. 2008](#), [Fernandes et al. 2014](#)). [Psaradellis and Sermpinis \(2016\)](#) extend these approaches by successfully adopting them in a hybrid genetic algorithm support vector regression of implied volatility index data (see, e.g., [Dunis et al. 2013](#), [Sermpinis et al. 2014](#), for further details). Besides the drawback of the studies being conducted at a daily frequency, another common disadvantage is that the discrete models employed ignore the underlying implied VIX evolution dynamics, that is, how the index moves from time $t - 1$ to time t . Therefore, the underlying stochastic process that generates intraday VIX observations cannot be determined. However, functional data analysis techniques, such as the techniques we implement in the present study, can be used to extract additional information beneath a time series of functions, assisting the investigation of essential sources of pattern and variation (see, e.g., [Horváth and Kokoszka 2012](#), [Kokoszka and Reimherr 2017](#), [Ramsay and Silverman 2005](#), Chapter 1).

Functional time series is attracting an ever-increasing focus, leading to a rapidly growing body of research on modelling and forecasting. From a parametric perspective, [Bosq \(2000\)](#) proposed the functional autoregressive of order 1 (FAR(1)) and derived one-step-ahead forecasts that are based on a regularised form of the Yule-Walker equations. FAR(1) was later extended to a FAR(p), under which order p can be determined via [Kokoszka and Reimherr \(2013\)](#)’s hypothesis testing procedure. [Aue et al. \(2015\)](#) introduced a vector autoregressive (VAR) model based on the functional principal component analysis (PCA), together with a forecasting method based on VAR forecasts of principal component scores. This VAR model can also be considered as an extension of [Hyndman and Shang \(2009\)](#), where principal component scores are forecast by a univariate time

series forecasting method, such as the autoregressive integrated moving average (ARIMA) method. [Klepsch and Klüppelberg \(2016\)](#) proposed the functional moving average (FMA) process with an innovations algorithm to obtain the best linear predictor. [Klepsch et al. \(2017\)](#) extended the FAR and FMA processes to a vector autoregressive moving average model, which can be considered a more straightforward estimation approach of the functional autoregressive moving average (FARMA) model. Recently, [Li et al. \(2018\)](#) considered long-range dependent curve time series and proposed a functional autoregressive fractionally integrated moving average model. From a nonparametric perspective, [Besse et al. \(2000\)](#) introduced functional kernel regression to model the temporal dependence by a similarity measure characterised by semi-metric, bandwidth, and kernel functions. From a semi-parametric perspective, [Aneiros-Pérez and Vieu \(2008\)](#) proposed a semi-functional partial linear model that combines parametric and nonparametric models, enabling consideration of additive covariates and the use of a continuous path in the past to predict future values of a stochastic process. Apart from the estimation of a conditional mean, [Hörmann et al. \(2013\)](#) consider a functional analogy of the autoregressive conditional heteroscedasticity model for modelling conditional variance. Recently, [Kokoszka et al. \(2017\)](#) proposed a portmanteau test for measuring autocorrelation under a functional analogue of the generalised autoregressive conditional heteroscedasticity model.

Among the various modelling and forecasting methods for functional time series, such as those listed above, many adopt functional PCA as a dimension reduction tool. Functional PCA can summarise an infinite-dimensional object into finite dimensions, sacrificing little information in the process. We utilise the functional PCA approach, also considered in [Hyndman and Shang \(2009\)](#) and [Aue et al. \(2015\)](#), to decompose a time series of functions into a set of functional principal components and their corresponding principal component scores. Correlation among each set of principal component scores obtained in functional PCA decomposition possesses temporal dependence information of the original functional time series. To forecast principal component scores, [Hyndman and Shang \(2009\)](#) considered a univariate time series forecasting method (e.g., autoregressive integrated moving average), while [Aue et al. \(2015\)](#) apply a multivariate time series forecasting method (e.g., a vector autoregressive model). Conditioning on the past curves and estimated functional principal components, point forecasts can be obtained by multiplying the forecast principal component scores by the estimated functional principal components. This type of algorithm is referred to as the 'TS method' in the remaining sections because it relies on either univariate or multivariate time series forecasting methods.

This paper is closely related to a previous study involving forecasting five-minute S&P 500 index returns using various functional time series methods with dynamic updating by [Shang \(2017\)](#). [Shang \(2017\)](#) noted that applying the forecasting methods to intraday financial data with denser time intervals (less than 5 minutes) would be a validation of the proposed forecasting

methods. We provide this validation through the use of the higher resolution 15-second VIX data. Our study contributes further by utilising robust approaches that control for the destabilising effect of outliers, and by extending the application of functional techniques to further the understanding of intraday implied VIX movements. We uncover intraday VIX patterns that allow academics and practitioners to better prepare for extreme events and given the proliferation of VIX-related futures, options, and exchange-traded products. While it is not our primary focus, our research could form the basis of idiosyncratic high frequency hedging and trading strategies.

The rest of this paper is structured as follows. In Section 2, we describe the VIX data employed in our study. In Section 3, we introduce the functional PCA, the robust functional PCA, and the robust regularised singular value decomposition. We present several dynamic updating methods for updating point forecasts in Section 4 and introduce some dynamic updating methods for updating interval forecasts in Section 5. Using the forecast error measures shown in Section 6, we examine point and interval forecast accuracies in Section 7. Conclusions are presented in Section 8, along with some ideas on how the present research can be extended.

2 Intraday CBOE volatility index

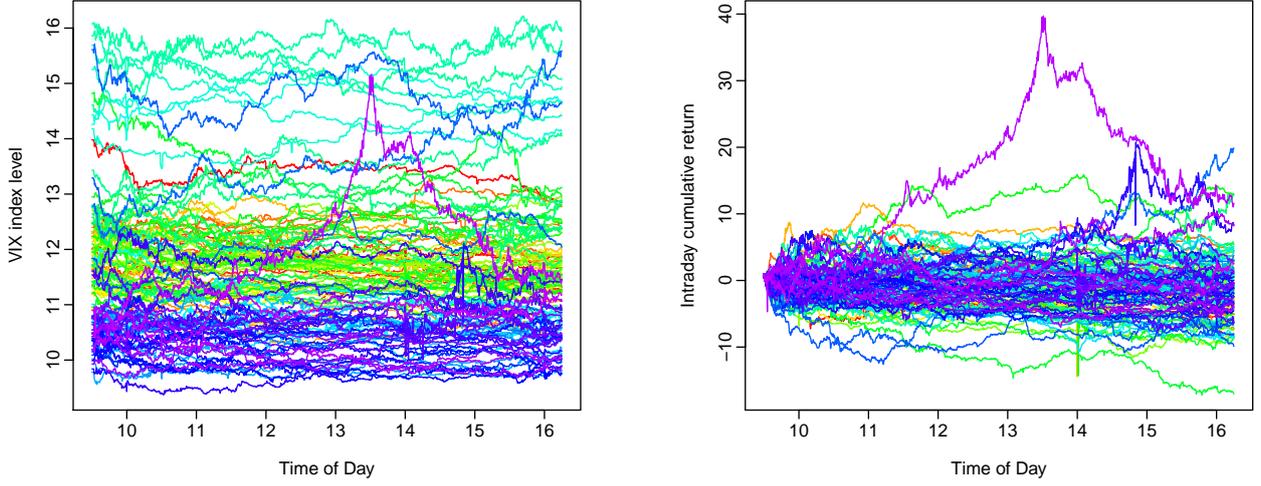
As the CBOE calculates the VIX index every 15 seconds¹, we adopt the highest resolution dataset and consider 15-second VIX index values from 3 January 2017 to 30 June 2017. Via the functional time series analysis techniques, non-stationary daily curves of intraday VIX values are transformed into daily curves of cumulative intraday returns (CIDRs, see, e.g., Gabrys et al. 2010). CIDRs have been proved to be stationary by Kokoszka et al. (2014). Although the VIX itself is not directly tradable, the approach is employed as it offers a transformation that facilitates the construction of stationary functional time series. Let $P_i(t_j)$ denote the daily VIX value at time t_j ($j = 1, \dots, m$) on day i ($i = 1, \dots, n$). CIDRs are computed as

$$R_i(t_j) = 100 \times [\ln P_i(t_j) - \ln P_i(t_1)], \quad (1)$$

where $\ln(\cdot)$ represents the natural logarithm and time intervals between t_{j-1} and t_j are 15 seconds in length. Figure 1 indicates that the curved shape of CIDRs is similar to the shape of the VIX curves at the original scale.

Thus, we consider constructing continuous curves of CIDRs on the observed VIX index values

¹The use of 15-second frequency is because it is the highest frequency available for the VIX index. Some of its derivatives may trade at a higher frequency, but the index value is only recalculated and released on a 15-second basis. So to provide us with the most current information and the highest level of granularity to model its evolution, we use these 15-second VIX quotes. This data frequency is particularly relevant for the dynamic functional time series approaches we propose.



(a) Functional time series curves of intraday VIX index

(b) Functional time series curves of CIDR VIX index

Figure 1: Graphical displays of VIX index functional data curves from 3 January 2017 to 30 June 2017.

as

$$\mathcal{X}_i(t) = R_i(t_j), \quad t \in (15(j-1), 15j], \quad \text{for } j = 1, \dots, \tau,$$

where τ denotes the total number of realised 15-second intraday cumulative returns. In this study, we consider VIX values reported by CBOE between 9:30:00 to 16:15:00 on each trading day. Therefore, τ has a maximal value of 1621. Once the functional time series is constructed, we work directly with this continuous series going forward.

To construct continuous CIDRs, we initially group 201,500 discrete observations obtained between 9:30:00 to 16:15:00 into 125 daily curves. We observe that on different trading days, timings of the first and the last VIX values can vary slightly. To ensure each daily curve is of the same length, we adopt the linear interpolation algorithm (Hyndman et al. 2018) in R software (R Core Team 2018) to fill in missing values (if any) at both ends of each daily curve. CIDRs are then computed based on Equation (1), leading to $N = 202,625$ discrete intraday cumulative returns separated into $n = 125$ CIDRs.

Traditionally, the presence of outliers can diminish the performance of functional time series models. Therefore, we showcase that our approach is robust to the presence of such outliers. According to Febrero et al. (2007) and Hyndman and Shang (2010), a functional outlier is a curve generated by a stochastic process that has a different distribution from that of standard curves. Based on this general definition, outliers can be further grouped into categories, such as magnitude outlier, shape outlier, and partial outlier (see Sguera et al. 2016, for more details). Using the functional outlier detection method proposed by Hyndman and Shang (2010), among the 125 curves the seven most volatile daily curves (about 5%)² are suspected to be outliers, as shown

²The proportion of outlying curves is specified to be around 5% which is an arbitrary but commonly used threshold in functional data analysis.

in Figure 2. Identified outliers that display abnormal shapes and magnitudes correspond to the dates 30 January 2017, 21 March 2017, 27 March 2017, 17 May 2017, 19 May 2017, 9 June 2017 and 29 June 2017 (as highlighted by the coloured lines in Figure 2b). These days correspond with events of special economic significance as we now outline. The first highlighted trading day of 30 January 2017 witnessed a stock market hit by US travel ban fears as President Trump placed a ban on immigration from seven Muslim-majority countries. The most volatile curve displayed corresponded to 29 June 2017 when technology stocks of S&P 500 slumped suddenly amid concerns of overvaluation and monetary policy tightening from the Federal Reserve causing investors to dump³.

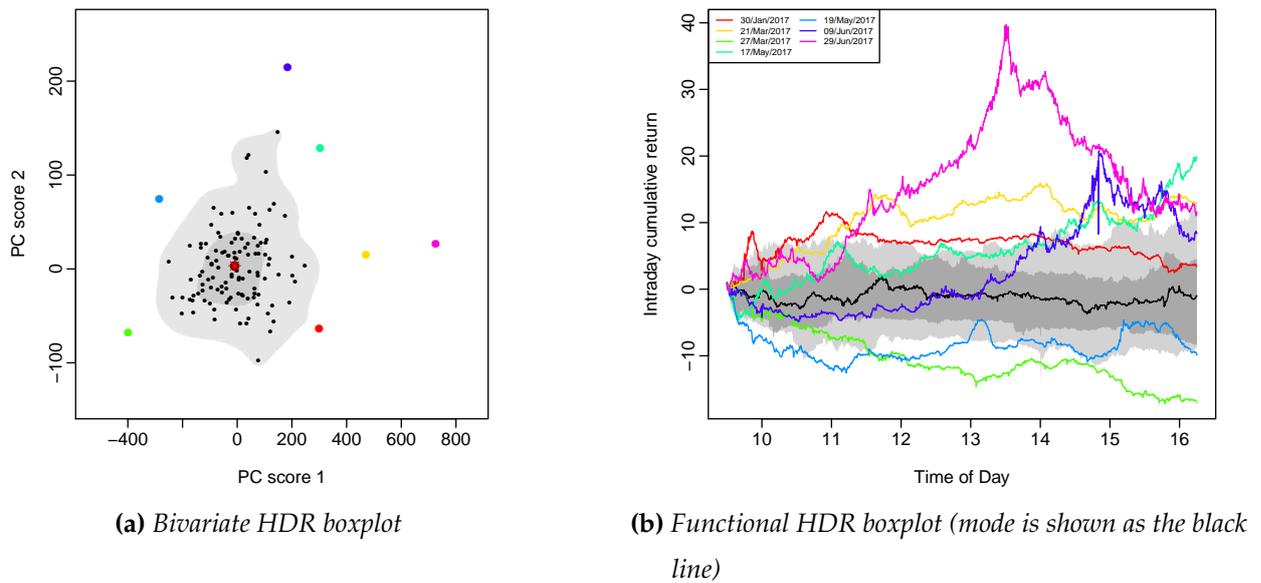


Figure 2: A functional outlier detection method, namely the highest density region, is adapted to identify five outliers representing about 5% of the total number of curves.

We consider functional time series forecasting methods that are robust to the effect of outliers. To exhibit this desirable property, we have not excluded any of those seven potentially outlying curves mentioned in this section from either our parameter estimation or forecasting processes.

3 Forecasting method

We employ the functional principal component regression of Hyndman and Shang (2009), Shen (2009), Hays et al. (2012) and Aue et al. (2015) to model and forecast functional time series of CIDRs. We adopt the functional PCA as the primary forecasting technique because it plays a vital role in the development of functional data analysis (for reviews, see Hall 2011, Morris 2015, Shang

³A comprehensive discussion of these special economic events leading to unusual market movements are included in our Supplement.

2014). To deal with outliers identified, we consider a robust functional PCA method (Hyndman and Ullah 2007) and a robust regularised singular value decomposition (Zhang et al. 2013).

3.1 Functional principal component regression

We denote an arbitrary stationary functional time series by $(\mathcal{X}_i : i \in \mathbb{Z})$. \mathcal{X}_i are assumed to be located in the Hilbert space, $\mathcal{H} = L^2(\mathcal{I})$ which has the inner product $\langle x, y \rangle = \int_{\mathcal{I}} x(t)y(t)dt$, where t is a continuum and $\mathcal{I} \in [1, p]$ represents a function support range. \mathcal{X}_i are assumed to be square integrable with finite squared norm $\|\mathcal{X}_i\|^2 = \int_{\mathcal{I}} \mathcal{X}_i^2(t)dt < \infty$. Using \mathcal{X} to denote a set of all \mathcal{X}_i defined on a common probability space (Ω, A, P) , we have $\mathcal{X} \in L^p_H(\Omega, A, P)$, indicating that for some positive value $p > 0$, $E(\|\mathcal{X}\|^p) < \infty$. When $p = 1$, it gives the mean curve of the stationary functional time series $\mathcal{X}(t)$, defined as $\mu(t) = E[\mathcal{X}(t)]$; when $p = 2$, it gives the covariance operator $\mathcal{K}(s, t) = \text{cov}[\mathcal{X}(s), \mathcal{X}(t)] = E\{[\mathcal{X}(s) - \mu(s)][\mathcal{X}(t) - \mu(t)]\}$. According to Karhunen (1946) and Loève (1946), the covariance operator can be defined as

$$\mathcal{K}(s, t) = \sum_{k=1}^{\infty} \lambda_k \phi_k(s) \phi_k(t), \quad s, t \in \mathcal{I},$$

where $\phi_k(t)$ denotes the k^{th} orthonormal principal component, and λ_k denotes the k^{th} eigenvalue.

Separability of the Hilbert space allows us to apply the Karhunen-Loève expansion to approximate the stochastic process \mathcal{X} based on the principal component score β_k given by the projection of $[\mathcal{X}(t) - \mu(t)]$ in the direction of the k^{th} eigenfunction ϕ_k as

$$\mathcal{X}(t) = \mu(t) + \sum_{k=1}^{\infty} \beta_k \phi_k(t). \quad (2)$$

Dimension reduction can be achieved by utilising Equation (2) and truncating the first K principal components. In this way, we assume that the K -dimensional vector $(\beta_1, \dots, \beta_K)$ contains most of the useful information in $\mathcal{X}(t)$, leading to an adequate approximation as

$$\mathcal{X}(t) = \mu(t) + \sum_{k=1}^K \beta_k \phi_k(t) + e(t), \quad (3)$$

where $e(t)$ represents model residuals after excluding the first K principal components. There are at least four approaches to selecting the number of retained principal components K : (1) the scree plot or the fraction of variance explained by the first several functional principal components (Chiou et al. 2012); (2) the pseudo-versions of the Akaike information criterion (AIC) and Bayesian information criterion (Yao et al. 2005); (3) the cross-validation with one-curve-leave-out (Rice and Silverman 1991); (4) the bootstrap technique (Hall and Vial 2006). We use the first method to find the value of K satisfying the following condition:

$$K = \underset{K:K \geq 1}{\text{argmin}} \left\{ \sum_{k=1}^K \hat{\lambda}_k / \sum_{k=1}^{\infty} \hat{\lambda}_k \mathbb{1}_{\{\hat{\lambda}_k > 0\}} \geq \delta \right\},$$

where $\delta = 90\%$ and $\mathbb{1}_{\{\hat{\lambda}_k > 0\}}$ to exclude possible zero eigenvalues⁴.

In practice, construction of a functional time series $\mathcal{X}(t) = \{\mathcal{X}_1(t), \dots, \mathcal{X}_n(t)\}$ utilises the empirical mean function $\hat{\mu}(t) = \frac{1}{n} \sum_{i=1}^n \mathcal{X}_i(t)$ and the empirical covariance function $\hat{\mathcal{K}}(s, t)$ estimated based on data. Estimated functional principal components functions $\Phi(t) = [\hat{\phi}_1(t), \dots, \hat{\phi}_k(t)]$ are then extracted from the empirical covariance function. The h -step-ahead point forecasts of curve $\mathcal{X}_{n+h}(t)$ can then be obtained through a conditional expectation as

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

where $\hat{\beta}_{n+h|n,k}$ represents the point forecasts of $\beta_{n+h,k}$. A univariate time series forecasting method and a multivariate time series forecasting method are used in the analysis of our VIX data and discussed in Section 3.4.

3.2 Robust functional principal component analysis

The presence of outliers often affects the estimation of functional principal components from the covariance operator $\mathcal{K}(s, t) = \text{cov}[\mathcal{X}(s), \mathcal{X}(t)]$ and leads to inferior estimation and forecast accuracies. Hyndman and Ullah (2007) introduced a two-step algorithm of extracting robust functional principal components by down-weighting the effect of outliers. This procedure begins by initially implementing the robust functional PCA algorithm as proposed by Hubert et al. (2002), calculating the integrated squared error of curve i as

$$v_i = \int_{t \in \mathcal{I}} \left[\mathcal{X}_i(t) - \sum_{k=1}^K \beta_{i,k} \phi_k(t) \right]^2 dt, \quad \text{for } i = 1, \dots, n.$$

By this algorithm, outlying curves tend to have large values of v_i . A set of weights can be assigned to curves satisfying $v_i < s + \lambda \sqrt{s}$, where a tuning parameter λ controls the amount of robustness and s represents the median of $\{v_1, \dots, v_n\}$. According to Hyndman and Ullah (2007), the efficiency of this procedure follows a cumulative normal distribution. We specify $\lambda = 2.33$ to ensure the efficiency equals to $\Phi(2.33/\sqrt{2}) = 95\%$, suggesting that 5% of curves are to be classified as outliers.

3.3 Robust regularised singular value decomposition

Apart from conducting robust PCA, we also consider the robust regularised singular value decomposition (RobRSVD) method of Zhang et al. (2013) for analysing functional data. RobRSVD

⁴Horváth and Kokoszka (2012) compared different methods of selecting the number of retained functional principal components and argued that the cumulative percentage of total variance method worked best throughout applications in their book. Hence, we have adopted this method in our study of VIX.

extends the idea of the traditional singular value decomposition (SVD), which aims to find a sequence of rank-one matrix approximations of a data matrix (Gabriel and Zamir 1979), to obtain a series of robust regularised rank-one matrix approximations. Implementation of the RobRSVD method involves computing the first pair of singular vectors via solving a least-squares problem as

$$(\mathbf{u}_1, \mathbf{v}_1) = \underset{\mathbf{u}, \mathbf{v}}{\operatorname{argmin}} \left\{ \rho \left(\mathcal{X} - \mathbf{u}\mathbf{v}^\top \right) + \mathcal{P}_\lambda(\mathbf{u}, \mathbf{v}) \right\},$$

where \mathbf{u} and \mathbf{v} are m -dimensional and n -dimensional vectors respectively, $\rho(\cdot)$ is a robust loss function, $\mathcal{P}_\lambda(\mathbf{u}, \mathbf{v})$ is a two-way roughness penalty to ensure smoothness for the \mathbf{u} and \mathbf{v} , and λ is a vector of penalty parameters. After removing the effects of the first pairs of \mathbf{u} and \mathbf{v} , following pairs of vectors can be obtained by applying the method to the model residuals from the previous step. Features of this sequential approach include allowing the different pairs of singular vectors to have different levels of smoothness. According to Zhang et al. (2013), the nonnegative, symmetric loss function for a rank-one approximation of the matrix \mathcal{X} is defined as

$$\rho \left(\frac{\mathcal{X} - \mathbf{u}\mathbf{v}^\top}{\sigma} \right) = \sum_{i=1}^m \sum_{j=1}^n \rho \left(\frac{\mathcal{X}_i(t_j) - u_i v_j}{\sigma} \right), \quad (4)$$

where σ is a scale parameter measuring the variability in the approximation error. In practice, the value of σ can be estimated by the conventional normalised median absolute deviation method, as defined by Maronna et al. (2006) as

$$\hat{\sigma} = \frac{1}{0.675} \operatorname{Med}_{ij} (|r_{ij}|, r_{ij} \neq 0),$$

where residuals are defined as $r_{ij} = \mathcal{X} - \hat{\mathbf{u}}\hat{\mathbf{v}}^\top$. Huang et al. (2009) suggested a specific form of the penalty function

$$\mathcal{P}_\lambda(\mathbf{u}, \mathbf{v}) = \lambda_u \mathbf{u}^\top \Omega_u \mathbf{u} \cdot \|\mathbf{v}\|^2 + \lambda_v \mathbf{v}^\top \Omega_v \mathbf{v} \cdot \|\mathbf{u}\|^2 + \lambda_u \mathbf{u}^\top \Omega_u \mathbf{u} \cdot \lambda_v \mathbf{v}^\top \Omega_v \mathbf{v},$$

where Ω_u and Ω_v are symmetric and nonnegative definite penalty matrices to the left and right singular vectors respectively, and $\|\cdot\|$ is the Euclidean norm. The implementation of Equation (4) includes a Huber's function defined as

$$\rho_\theta(x) = \begin{cases} x^2 & \text{if } |x| \leq \theta \\ 2\theta|x| - \theta^2 & \text{if } |x| > \theta \end{cases},$$

where a smaller value of the smoothness parameter θ gives a more robust estimation. A commonly used $\theta = 1.345$ is adopted because it produces 95% efficiency for normal errors (Huber and Ronchetti 2009).

3.4 Univariate and multivariate time series forecasting methods

We consider the univariate time series forecasting method of Hyndman and Shang (2009) which utilises an ARIMA model. This procedure can be applied to non-stationary time series that contains

a stochastic trend component. The VIX data considered spans 125 days without any apparent seasonality. Therefore, an ARIMA model is defined in the following general form:

$$(1 - \phi_1 \mathcal{B} - \dots - \phi_p \mathcal{B}^p) (1 - \mathcal{B})^d \beta_k = \gamma + (1 + \theta_1 \mathcal{B} + \dots + \theta_q \mathcal{B}^q) \mathbf{w}_k,$$

where γ represents the intercept, (ϕ_1, \dots, ϕ_p) denote the coefficients associated with the autoregressive component, $(\theta_1, \dots, \theta_q)$ denote the coefficients associated with the moving average component, \mathcal{B} denotes the backshift operator, d denotes the differencing operator, $\beta_k = (\beta_{1,k}, \dots, \beta_{n,k})$ represents the k^{th} estimated principal component scores, and $\mathbf{w}_k = (w_{1,k}, \dots, w_{n,k})$ represents the error term. We select the optimal ARIMA model based on an information criterion, and subsequently estimate parameters in the optimal model via the maximum likelihood estimation method. **In practice, we use the automatic ARIMA algorithm of Hyndman and Khandakar (2008) to select the optimal model based on the corrected AIC (an information criterion that performs well for small sample sizes).**

For data that have strict diagonal auto-covariances at all lags, univariate time series forecasting methods can be accurate and time efficient (see, e.g. [Aue et al. 2015](#)). Given the vectors of functional principal component scores may have the dependence structure of the principal component score matrix, a univariate functional time series modelling process may result in a loss of information. [Aue et al. \(2015\)](#) introduced a multivariate time series model to counter this problem. The most commonly used multivariate time series model is the vector autoregressive (VAR) model, for the following reasons: (1) estimation of parameters for the VAR model is generally straightforward as ordinary least squares (OLS), maximum likelihood and Bayesian methods are all possible candidate procedures; (2) the properties of the VAR model have been studied (see, e.g. [Tsay 2013](#)); (3) the VAR model has a clear structure and can be considered as a type of multivariate multiple linear regression. To utilise the multivariate time series forecasting method, we define a VAR model of order ϑ if the multivariate principal component scores $\beta_k = [\beta_{1,k}, \dots, \beta_{n,k}]^\top$ satisfy

$$\beta_k = \phi_0 + \sum_{v=1}^{\vartheta} \phi_v \beta_{k-v} + \mathbf{a}_k,$$

where ϕ_0 is an n -dimensional constant vector, and ϕ_v are $n \times n$ matrices for $v > 0$ and $\phi_\vartheta \neq 0$ and \mathbf{a}_k is a set of independent and identically distributed (i.i.d.) random error vectors with a mean of zero and a positive-definite covariance matrix that has only positive eigenvalues. We use the corrected AIC to select the optimal order of ϑ .

Comparisons of point and interval forecast accuracies between the univariate ARIMA and multivariate VAR models can be found in the literature. For example, [Peña and Sánchez \(2007\)](#) noted that if the series (i.e., principal component score vectors) are very weakly related, considering the joint dynamics of the series can marginally improve the univariate time series forecasts. Conversely, when historical observations of one time series can significantly influence another

time series, the multivariate time series forecasts have many advantages compared to the univariate time series forecasts. We present a comparison between point and interval forecasts obtained by the univariate and multivariate time series forecasting methods in Section 7.

4 Updating point forecasts

Constructing functional time series based on segments of a longer univariate time series often encounters the problem that newly arriving curve data may not possess all the important features of the complete curve. After observing the first m_0 time periods on day $n + 1$ of $\mathcal{X}_{n+1}(t)$, denoted by $\mathcal{X}_{n+1}(t_e) = [\mathcal{X}_{n+1}(t_1), \dots, \mathcal{X}_{n+1}(t_{m_0})]^\top$, our interest is forecasting the data in the remaining part of the same curve, denoted by $\mathcal{X}_{n+1}^l(t)$, where $t \in \mathcal{I}_l$. As described in Section 3, the one-step-ahead TS forecast of

$$\hat{\mathcal{X}}_{n+1|n}^{l,TS}(t) = \hat{\mu}^l(t) + \sum_{k=1}^K \hat{\beta}_{n+1|n,k}^{TS} \hat{\phi}_k^l(t), \quad \text{for } l \in (m_0, p], \quad (5)$$

where $\hat{\mu}^l(t)$ represents the mean curve for the remaining time period and $\hat{\phi}_k^l(t)$ represents the k^{th} functional principal component of the remaining time period.

However, Equation (5) does not utilise the partially observed curve. For improving point forecast accuracy, dynamic updating methods are often considered because such methods incorporate newly observed information into forecasting the remaining curve of the day $n + 1$ (see also Shang and Hyndman 2011). In this section, several dynamic updating methods are introduced and compared in consideration of their ability to improve forecast accuracy.

4.1 Block moving (BM)

The BM method adopts the general structure of the TS method with redefined starting and ending time points for our CIDR curves. Treating time as a continuous variable, the BM method changes the support range of Equation (3) from $[1, p]$ to $(m_0, p] \cup [1, m_0]$. The consequences include a loss of data at the beginning of the first curve and complete subsequent partially observed curves. The data loss in the first curve will only have a minimal negative influence on the forecasts given enough data (i.e., a large number of curves) because the forecasts depend very little on observations from the distant past.

4.2 Ordinary least squares (OLS)

A regression approach based on the functional principal components obtained in Equation (3) can also be used to approximate the remaining part of the most recent curve. Let \mathcal{F}_e be the $m_0 \times K$ matrix of the response variable, whose $(j, k)^{\text{th}}$ entry is $\hat{\phi}_k(t_j)$, for $1 \leq j \leq m_0$ and $1 \leq k \leq K$. Let $\mathbf{X}_{n+1}(t_e)$ be the $m_0 \times 1$ matrix of the response variable, $\beta_{n+1} = [\beta_{n+1,1}, \dots, \beta_{n+1,K}]^\top$ and

$\varepsilon_{n+1}(t_e) = [\varepsilon_{n+1}(t_1), \dots, \varepsilon_{n+1}(t_{m_0})]^\top$. The regression model used to estimate coefficients β_{n+1} has the following equation

$$\mathbf{X}_{n+1}^*(t_e) = \mathcal{F}_e \beta_{n+1} + \varepsilon_{n+1}(t_e),$$

where $\mathbf{X}_{n+1}^*(t_e) = \mathbf{X}_{n+1}(t_e) - \hat{\mu}(t_e)$ denotes the mean-adjusted response variable. Using the OLS estimation method, we have

$$\hat{\beta}_{n+1}^{\text{OLS}} = (\mathcal{F}_e^\top \mathcal{F}_e)^{-1} \mathcal{F}_e^\top \mathbf{X}_{n+1}^*(t_e). \quad (6)$$

Then, the OLS forecasts of $\mathcal{X}_{n+1}^l(t)$ at discretised time points are given by

$$\hat{\mathbf{X}}_{n+1}^{\text{OLS}}(t_l) = \hat{\mu}(t_l) + \sum_{k=1}^K \hat{\beta}_{n+1,k}^{\text{OLS}} \hat{\phi}_k(t_l),$$

where $t_l = \{t_{m_0+1}, \dots, t_\tau; \tau \leq q\}$ represents the discretised time points in the remaining period⁵.

4.3 Ridge regression (RR)

To conduct the inverse computation in Equation (6) in practice, a sufficiently large number of observations for $\hat{\beta}_{n+1}^{\text{OLS}} = [\hat{\beta}_{1,n+1}^{\text{OLS}}, \dots, \hat{\beta}_{K,n+1}^{\text{OLS}}]^\top$ are necessary. To address this problem, we consider the ridge regression method of Hoerl and Kennard (1970), using the functional principal components as predictors and the partially observed data as responses. The RR method has a feature of shrinking the regression coefficient estimates towards zero, achieved by minimising a penalised residual sum of squares:

$$\text{argmin} \left\{ (\mathbf{X}_{n+1}^*(t_e) - \mathcal{F}_e \beta_{n+1})^\top (\mathbf{X}_{n+1}^*(t_e) - \mathcal{F}_e \beta_{n+1}) + \lambda \beta_{n+1}^\top \beta_{n+1} \right\},$$

where $\lambda > 0$ is a parameter that controls the amount of shrinkage. Taking the first derivative with respect to β_{n+1} leads to the estimated coefficients as

$$\hat{\beta}_{n+1}^{\text{RR}} = \left(\mathcal{F}_e^\top \mathcal{F}_e + \lambda \mathbf{I}_K \right)^{-1} \mathcal{F}_e^\top \mathbf{X}_{n+1}^*(t_e),$$

where \mathbf{I}_K is a $(K \times K)$ identity matrix. The $\hat{\beta}_{n+1}^{\text{RR}}$ has the following properties: (i) when $0 < \lambda < \infty$, $\hat{\beta}_{n+1}^{\text{RR}}$ is a weighted average between 0 and $\hat{\beta}_{n+1}^{\text{OLS}}$; (ii) when $\lambda \rightarrow 0$, $\hat{\beta}_{n+1}^{\text{RR}}$ approaches to $\hat{\beta}_{n+1}^{\text{OLS}}$; (iii) when $\lambda \rightarrow \infty$, $\hat{\beta}_{n+1}^{\text{RR}}$ goes to 0. With an optimal selection of λ , the RR forecasts of $\mathcal{X}_{n+1}^l(t)$ at discretised time points are given by

$$\hat{\mathbf{X}}_{n+1}^{\text{RR}}(t_l) = \hat{\mu}(t_l) + \sum_{k=1}^K \hat{\beta}_{n+1,k}^{\text{RR}} \hat{\phi}_k(t_l).$$

⁵When updating point forecasts for a particular day, we keep using the same K obtained in the model estimation step as per Equation (3).

4.4 Penalised least squares (PLS)

The RR method avoids the singularity problem associated with $\widehat{\beta}_{n+1}^{\text{OLS}}$, but it does not use the TS forecasts $\widehat{\beta}_{n+1|n}^{\text{TS}}$. The PLS method shrinks regression coefficient estimates towards $\widehat{\beta}_{n+1|n}^{\text{TS}}$ by minimising a penalised residual sum of squares:

$$\operatorname{argmin} \left\{ (\mathbf{X}_{n+1}^*(t_e) - \mathcal{F}_e \beta_{n+1})^\top (\mathbf{X}_{n+1}^*(t_e) - \mathcal{F}_e \beta_{n+1}) + \lambda (\beta_{n+1} - \widehat{\beta}_{n+1|n}^{\text{TS}})^\top (\beta_{n+1} - \widehat{\beta}_{n+1|n}^{\text{TS}}) \right\}. \quad (7)$$

In Equation (7), the first term measures the ‘‘goodness of fit’’ and the second term penalises the departure of the regression coefficient estimates from the $\widehat{\beta}_{n+1|n}^{\text{TS}}$. Taking the first derivative with respect to β_{n+1} leads to the estimated coefficients as

$$\widehat{\beta}_{n+1}^{\text{PLS}} = \left(\mathcal{F}_e^\top \mathcal{F}_e + \lambda \mathbf{I}_K \right)^{-1} \left(\mathcal{F}_e^\top \mathbf{X}_{n+1}^*(t_e) + \lambda \widehat{\beta}_{n+1|n}^{\text{TS}} \right). \quad (8)$$

The $\widehat{\beta}_{n+1}^{\text{PLS}}$ has the following properties: (i) when $0 < \lambda < \infty$, $\widehat{\beta}_{n+1}^{\text{PLS}}$ is a weighted average between $\widehat{\beta}_{n+1|n}^{\text{PLS}}$ and $\widehat{\beta}_{n+1}^{\text{OLS}}$; (ii) when $\lambda \rightarrow 0$, $\widehat{\beta}_{n+1}^{\text{PLS}}$ approaches to $\widehat{\beta}_{n+1}^{\text{OLS}}$ given $\left(\mathcal{F}_e^\top \mathcal{F}_e \right)^{-1}$ exists; (iii) when $\lambda \rightarrow \infty$, $\widehat{\beta}_{n+1}^{\text{PLS}}$ go to $\widehat{\beta}_{n+1}^{\text{OLS}}$. With an optimal selection of λ , the PLS forecasts of $\mathcal{X}_{n+1}^l(t)$ at discretised time points are given by

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

4.5 Selection of shrinkage parameters in the RR and PLS methods

Both the RR method and the PLS method require the optimal selection of λ . To minimise forecast errors, the choice of λ should also vary as we observe an increasing number of data points in a trading day. We adopt a holdout forecast evaluation method to select λ . We divide a time series of functions into a training sample (including days from 1 to 85) and a testing sample (including days from 86 to 125). We further divide the training sample into a training set (including days from 1 to 43) and a validation set (including days from 44 to 85). The optimal values of λ for different updating periods are determined by minimising the averaged forecast error criteria within the validation set (see Section 6).

4.6 Functional linear regression

The OLS, RR, and PLS methods share a common feature of considering a discretised data approach to produce updates for the remaining period. To take advantage of functional data analysis, which allows observations to be considered as realisations of a continuous function, we also consider the functional linear regression of [Chiou et al. \(2012\)](#) (see also [Müller et al. 2011](#)). For a detailed overview of functional linear regression, refer to [Müller \(2005\)](#) and [Morris \(2015\)](#). The functional linear regression can be expressed as

$$\mathcal{X}_{n+1}^l(t) = \mu^l(t) + \int_{s \in \mathcal{I}_e} [\mathcal{X}_{n+1}^e(s) - \mu^e(s)] \beta(s, t) ds + e_{n+1}^l(t) \quad s \in \mathcal{I}_e, t \in \mathcal{I}_l, \quad (9)$$

where $\mathcal{I}_e \in [1, m_0]$ and $\mathcal{I}_l \in (m_0, p]$ represent two function support ranges for the partially observed and remaining segments of curve on day $n + 1$, respectively; $\mu^e(s)$ and $\mu^l(t)$ represent two mean functions for the partially observed data and remaining data periods, respectively; $\mathcal{X}_{n+1}^e(s)$ and $\mathcal{X}_{n+1}^l(t)$ represent functional predictor and functional response variables, respectively. Equation (9) can be viewed as function-on-function linear regression (see also Ramsay and Silverman 2005, Chapter 16), where $\beta(s, t)$ and $e_{n+1}^l(t)$ denote the regression coefficient function and error function, respectively.

Estimation of the regression coefficient function $\beta(s, t)$ involves projecting a time series of functions onto functional principal component scores. Using functional PCA, we obtain

$$\begin{aligned}\mathcal{X}_i^e(t) &= \mu^e(t) + \sum_{k=1}^{\infty} \zeta_{i,k} \phi_k^e(t) & \mathcal{X}_i^l(t) &= \mu^l(t) + \sum_{m=1}^{\infty} \zeta_{i,m} \Psi_m^l(t) \\ &= \mu^e(t) + \sum_{k=1}^K \zeta_{i,k} \phi_k^e(t) + \eta_i^e(t), & &= \mu^l(t) + \sum_{m=1}^M \zeta_{i,m} \Psi_m^l(t) + v_i^l(t),\end{aligned}$$

where $\phi_k^e(t)$ and $\Psi_m^l(t)$ represent the k^{th} and m^{th} functional principal components for the partially observed and remaining segments of curve on day $n + 1$, respectively; $\zeta_{i,k}$ and $\zeta_{i,m}$ are the principal component scores of $\mathcal{X}_i^e(t)$ and $\mathcal{X}_i^l(t)$, respectively; K and M denote the number of retained components, respectively; $\eta_i^e(t)$ and $v_i^l(t)$ represent the error functions associated with the partially observed data and remaining data periods, respectively, due to finite truncations.

Let $\zeta_m = [\zeta_{1,m}, \dots, \zeta_{n,m}]^\top$ and $\xi_k = [\zeta_{1,k}, \dots, \zeta_{n,k}]^\top$. Using $\zeta = [\zeta_1, \dots, \zeta_M]$ as the response variable and $\xi = [\xi_1, \dots, \xi_K]$ as a predictor, the equation considered can be expressed as

$$\zeta = \xi \times \varsigma, \quad (10)$$

where ς can be estimated by OLS, leading to

$$\hat{\varsigma} = \left(\xi^\top \xi \right)^{-1} \xi^\top \zeta, \quad (11)$$

with $\xi^\top \zeta$ jointly estimated by their cross-covariance structure

$$\int_{t \in \mathcal{I}} \int_{s \in \mathcal{I}_e} \phi_k(s) \Psi_m(t) \text{cov} \left[\mathbf{X}^e(s), \mathbf{X}^l(t) \right] ds dt, \quad k = 1, \dots, K, \quad m = 1, \dots, M,$$

where $\mathbf{X}^e(s) = [\mathcal{X}_1^e(s), \dots, \mathcal{X}_n^e(s)]^\top$ and $\mathbf{X}^l(t) = [\mathcal{X}_1^l(t), \dots, \mathcal{X}_n^l(t)]^\top$ represent two vectors of functions that correspond to the partially observed and remaining segments of curve on day $n + 1$, respectively.

From Equation (9), we can obtain a point forecast of $\mathcal{X}_{n+1}^l(t)$ as

$$\hat{\mathcal{X}}_{n+1}^l(t) = \hat{\mu}^l(t) + \sum_{m=1}^{\infty} \zeta_{n+1,m} \Psi_m^l(t), \quad (12)$$

which has an approximation based on Equation (10) and Equation (12) as

$$\hat{\mathcal{X}}_{n+1}^l(t) \approx \hat{\mu}^l(t) + \hat{\xi}_{n+1} \times \hat{\varsigma} \times \hat{\Psi}^l(t),$$

where $\hat{\varsigma}$ is estimated from Equation (11), and $\hat{\Psi}^l(t) = [\hat{\Psi}_1^l(t), \dots, \hat{\Psi}_M^l(t)]$.

5 Interval forecast methods

Prediction intervals are often used to assess the probabilistic uncertainty associated with point forecasts. As emphasised in [Chatfield \(1993\)](#), it is essential to provide interval forecasts for the following reasons: (1) assess future uncertainty; (2) enable different strategies to be planned for a range of possible outcomes indicated by the interval forecasts; (3) compare forecasts from different methods more thoroughly; (4) explore different scenarios based on different assumptions. It is always essential to be clear about the sources of errors before quantifying forecast uncertainty via computations. In our functional principal component regression, we identify two sources of errors: (1) the errors in estimating the regression coefficient estimates; (2) the errors remain in the model residuals. In [Section 5.1](#), we describe a nonparametric bootstrap method for constructing one-step-ahead prediction intervals for the TS method, and in [Section 5.2](#) we demonstrate how the prediction intervals can be updated from the BM method, PLS method, and functional linear regression.

5.1 Nonparametric prediction interval

The focus of our study is on short-term time series forecasting; we define the one-step-ahead in-sample forecast errors for estimated principal component scores as

$$\widehat{\Omega}_{j,k} = \widehat{\beta}_{n-j+1,k} - \widehat{\beta}_{n-j+1|n-j,k}, \quad \text{for } j = 1, \dots, n - K,$$

where K represents the number of retained principal components in [Equation \(3\)](#). Sampling with replacement within these one-step-ahead forecast errors gives a bootstrap sample of $\beta_{n+h,k}$, denoted by $\widehat{\beta}_{n+1|n,k}^b$. For $h = 1$, we have

$$\widehat{\beta}_{n+1|n,k}^b = \widehat{\beta}_{n+1|n,k} + \widehat{\Omega}_{*,k}^b, \quad \text{for } b = 1, \dots, B,$$

where $\widehat{\Omega}_{*,k}^b$ denotes the bootstrapped forecast errors obtained by sampling with replacement from $(\widehat{\Omega}_{1,k}, \dots, \widehat{\Omega}_{n-K,k})$, and $B = 1000$ symbolises the number of bootstrap replications.

We assume the model residuals to be i.i.d random noise given the first K functional principal component decomposition in [Equation \(3\)](#) approximates data relatively well. We can then bootstrap the model residual function $\widehat{e}_{n+1}(t)$ by randomly sampling with replacement from the historical residual functions $\{\widehat{e}_1(t), \dots, \widehat{e}_n(t)\}$. Putting the two sources of errors together yields B bootstrapped forecasts of $\mathcal{X}_{n+1}(t)$, as defined by

$$\widehat{\mathcal{X}}_{n+1|n}^b(t) = \widehat{\mu}(t) + \sum_{k=1}^K \widehat{\beta}_{n+1|n,k}^b \widehat{\phi}_k(t) + \widehat{e}_{n+1}^b(t).$$

Therefore, $100(1 - \alpha)\%$ prediction intervals can be defined as $\alpha/2$ and $(1 - \alpha/2)$ empirical quantiles of $\{\widehat{\mathcal{X}}_{n+1|n}^1(t), \dots, \widehat{\mathcal{X}}_{n+1|n}^B(t)\}$. With a modification of function support range, this nonparametric prediction interval approach also works for the BM method.

5.2 Updating prediction interval

5.2.1 PLS method

We can dynamically update prediction intervals using a nonparametric bootstrap method with sequentially observed new data. The initial step involves bootstrapping B samples of the TS forecast regression coefficient function, denoting them by $\widehat{\beta}_{n+1|n}^{b,TS} = [\widehat{\beta}_{n+1|n,1'}^{b,TS}, \dots, \widehat{\beta}_{n+1|n,K}^{b,TS}]^\top$. The bootstrapped TS regression coefficient estimates will consequently lead to $\widehat{\beta}_{n+1}^{b,PLS}$ according to Equation (8). Then we obtain B replications of $\widehat{\mathbf{X}}_{n+1}^{b,PLS}(t_l)$ as

$$\widehat{\mathbf{X}}_{n+1}^{b,PLS}(t_l) = \widehat{\boldsymbol{\mu}}(t_l) + \sum_{k=1}^K \widehat{\beta}_{n+1,k}^{b,PLS} \widehat{\phi}_k(t_l) + \widehat{e}_{n+1}^b(t_l). \quad (13)$$

Therefore, $100(1 - \alpha)\%$ prediction intervals for the updated forecasts can be defined as $\alpha/2$ and $(1 - \alpha/2)$ empirical quantiles of $\{\widehat{\mathbf{X}}_{n+1}^{1,PLS}(t_l), \dots, \widehat{\mathbf{X}}_{n+1}^{B,PLS}(t_l)\}$.

5.2.2 Functional linear regression

Apart from Equation (13), which constructs the pointwise prediction intervals for the remaining segments of the curve on day $n + 1$, we can also use functional linear regression together with a bootstrap method to build pointwise prediction intervals. Utilising the functional linear regression in Equation (9), we can obtain the bootstrapped forecasts $\mathcal{X}_{n+1}^{l,b}(t)$ as

$$\widehat{\mathcal{X}}_{n+1}^{l,b}(t) = \mu^l(t) + \int_{s \in \mathcal{I}_e} [\mathcal{X}_{n+1}^e(s) - \mu(s)] \widehat{\beta}^b(s, t) ds + \widehat{e}_{n+1}^{l,b}(t), \quad s \in \mathcal{I}_e, \quad t \in \mathcal{I}_l, \quad (14)$$

where $\widehat{\beta}^b(s, t)$ represents the bootstrapped regression coefficient estimates, and $\widehat{e}_{n+1}^{l,b}(t)$ stands for the bootstrapped error function associated with the remaining time period. Parameter variation in the estimation of regression coefficients is captured by $\widehat{\beta}^b(s, t)$, while $\widehat{e}_{n+1}^{l,b}(t)$ measures the model variability in the fitted model.

We assume that the one-step-ahead forecast errors do not correlate and implement the i.i.d. bootstrap method by sampling with replacement from historical errors $\{\widehat{e}_K^l(t), \dots, \widehat{e}_n^l(t)\}$. We also use bootstrapping to obtain $\widehat{\beta}^b(s, t)$ in Equation (14) by sampling with replacement from the original functional time series expressed as

$$\mathcal{X}_i^b(t) = \mu(t) + \sum_{k=1}^{\min(n, \infty)} \beta_{i,k}^b \phi_k(t), \quad i = 1, \dots, n,$$

where $(\beta_{1,k}^b, \dots, \beta_{n,k}^b)$ denote the k^{th} bootstrapped principal component scores via maximum entropy (see also Shang 2018). With a set of bootstrapped data $\{\mathcal{X}_1^b(t), \dots, \mathcal{X}_n^b(t)\}$ available, we apply the functional linear regression of Equation (9) to obtain bootstrapped estimates of regression coefficient function.

For bootstrapping the multivariate time series $\{\beta_1, \dots, \beta_n\}$ we apply the maximum entropy bootstrap method (Vinod 2004) which has the following advantages: (1) stationarity of principal

component scores is not required; (2) time series ranks are computed by the method; (3) the bootstrap samples satisfy the ergodic theorem, central limit theorem, and mean preserving constraint; (4) the method can be applied to panel time series in which cross-covariance of the original data matrix is well preserved. In practice, we utilise the maximum entropy bootstrap algorithm described in [Vinod et al. \(2009\)](#), with implementation via the `meboot.pdata.frame` function of the `meboot` package in R ([R Core Team 2018](#)).

6 Measures of point forecast accuracy

6.1 Absolute and squared forecast errors

We calculate the point forecasts between the proposed methods and evaluate their forecast accuracy using mean absolute forecast error (MAFE) and mean squared forecast error (MSFE). The errors assess the absolute and squared differences between the forecasts and the actual values of the variable being forecast, with equations expressed as

$$\text{MAFE}_j = \frac{1}{q} \sum_{i=1}^q \left| \mathcal{X}_i(t_j) - \hat{\mathcal{X}}_i(t_j) \right|, \quad (15)$$

$$\text{MSFE}_j = \frac{1}{q} \sum_{i=1}^q \left[\mathcal{X}_i(t_j) - \hat{\mathcal{X}}_i(t_j) \right]^2, \quad (16)$$

where q represents the number of curves in the forecasting period, $\mathcal{X}_i(t_j)$ represents the actual holdout sample for the j^{th} time period in the i^{th} curve, and $\hat{\mathcal{X}}_i(t_j)$ represents the point forecasts for the holdout sample. To apply the updating algorithm detailed in Section 4, at least one data point on the i^{th} curve has to be observed. Using TS, BM, PLS and BM methods the updating period can start at the $j = 2^{\text{nd}}, \dots, 1620^{\text{th}}$ 15-second tick interval for the CIDRs of VIX data. However, the functional linear regression method requires the updating period to start from the $j = 6^{\text{th}}$ 15-second tick interval as the method is not suitable with initial zeros (added by linear interpolation, see Section 2 for more details) on each curve.

6.2 Mixed errors

We calculate mean mixed error (MME) functions that are asymmetric loss functions penalising incorrect predictions. [Brailsford and Faff \(1996\)](#), [Fuertes et al. \(2009\)](#) and [Kearney et al. \(2018\)](#) have previously considered by these measures in the evaluation of volatility forecasting methods. They are essential for option market participants as sellers (buyers) are more likely to worry about under (over)-prediction than buyers (sellers). The MME(U) is defined to penalise under-predictions more

heavily while the MME(O) gives more substantial penalties to under-predictions.

$$\text{MME}_j(U) = \frac{1}{q} \left[\sum_{i=i_1^O}^{i_N^O} |\mathcal{X}_i(t_j) - \hat{\mathcal{X}}_i(t_j)| + \sum_{i=i_1^U}^{i_N^U} \sqrt{|\mathcal{X}_i(t_j) - \hat{\mathcal{X}}_i(t_j)|} \right]$$

and

$$\text{MME}_j(O) = \frac{1}{q} \left[\sum_{i=i_1^O}^{i_N^O} \sqrt{|\mathcal{X}_i(t_j) - \hat{\mathcal{X}}_i(t_j)|} + \sum_{i=i_1^U}^{i_N^U} |\mathcal{X}_i(t_j) - \hat{\mathcal{X}}_i(t_j)| \right],$$

where q represents the number of curves in the forecasting period, $\{i_1^U, \dots, i_N^U; i_N^U \leq q\}$ stands for the under-predictions, and $\{i_1^O, \dots, i_N^O; i_N^O \leq q\}$ stands for the over-predictions.

6.3 Correct predictor of the direction of change

To measure how well our models can forecast the direction of movement, we also calculate the mean correct predictor of the direction of change (MCPDC) as the percentage of $\hat{\mathcal{X}}_i(t_j)$ that has the same sign as the corresponding observations $\mathcal{X}_i(t_j)$. MCPDC ignores the magnitude of errors and is also employed in other studies of the CBOE volatility index, for example in Bernales and Guidolin (2014) and Kearney et al. (2018).

6.4 Interval scores

We calculate the interval score of Gneiting and Raftery (2007) (see also Gneiting and Katzfuss 2014) to evaluate interval forecast accuracy. The one-step-ahead prediction intervals for each year in the forecasting period can be calculated at the $(1 - \alpha) \times 100\%$ nominal coverage probability. We consider the common case of the symmetric $(1 - \alpha) \times 100\%$ prediction interval, with lower and upper bounds serving as predictive quantiles at $\alpha/2$ and $1 - \alpha/2$, denoted by $\hat{\mathcal{X}}^l(t_j)$ and $\hat{\mathcal{X}}^u(t_j)$. The scoring rule of Gneiting and Raftery (2007) for the interval forecast at time point t_j is

$$S_\alpha \left[\hat{\mathcal{X}}^l(t_j), \hat{\mathcal{X}}^u(t_j); \mathcal{X}(t_j) \right] = \left[\hat{\mathcal{X}}^u(t_j) - \hat{\mathcal{X}}^l(t_j) \right] + \frac{2}{\alpha} \left[\hat{\mathcal{X}}^l(t_j) - \mathcal{X}(t_j) \right] \mathbb{1}\{\mathcal{X}(t_j) < \hat{\mathcal{X}}^l(t_j)\} \\ + \frac{2}{\alpha} \left[\mathcal{X}(t_j) - \hat{\mathcal{X}}^u(t_j) \right] \mathbb{1}\{\mathcal{X}(t_j) > \hat{\mathcal{X}}^u(t_j)\},$$

where $\mathbb{1}\{\cdot\}$ represents the binary indicator function, and α (customarily $\alpha = 0.2$) denotes the level of significance. The interval score rewards a narrow prediction interval, if and only if the true observation lies within the prediction interval. When $\mathcal{X}(t_j)$ is contained in the interval, and the distance between $\hat{\mathcal{X}}^l(t_j)$ and $\hat{\mathcal{X}}^u(t_j)$ is minimised, the optimal interval score will be achieved.

The mean interval score for each time point j is calculated by averaging interval scores over different days in the forecasting period as

$$\bar{S}_{\alpha,j} = \frac{1}{q} \sum_{i=1}^q S_\alpha \left[\hat{\mathcal{X}}_i^l(t_j), \hat{\mathcal{X}}_i^u(t_j); \mathcal{X}_i(t_j) \right], \quad j = 2, \dots, 1620,$$

where $S_\alpha[\hat{\mathcal{X}}_i^l(t_j), \hat{\mathcal{X}}_i^u(t_j); \mathcal{X}_i(t_j)]$ denotes the interval score for the i^{th} day of the forecasting period.

7 Results

We calculate and compare the point and interval forecast accuracies of the functional principal component regression with univariate and multivariate time series forecasting methods in Table 1. Generally, the multivariate (VAR) time series forecasting method produces slightly larger forecast errors than the univariate time series (ARIMA) forecasting method. This result may be due to there being more parameters to be estimated in the VAR model as opposed to the ARIMA model.

Table 1: Comparison of point and interval forecast accuracies between a univariate time series forecasting method and a multivariate time series forecasting method, for VIX CIDRs over the 40 days in the forecasting period.

	MSFE		MAFE		Interval score	
	ARIMA	VAR	ARIMA	VAR	ARIMA	VAR
Min	0.7939	1.3486	0.6900	0.9866	8.2059	8.7041
Median	8.1868	8.6180	2.3820	2.5446	9.1384	9.2167
Mean	22.0173	25.3357	3.0447	3.3261	15.7863	17.3692
Max	354.7139	342.1089	16.1869	16.0261	128.4479	124.2109

Hereafter, the ARIMA forecasting method is used to compare point and interval forecast accuracies between the standard functional PCA, the robust functional PCA, and the RobRSVD estimation for the TS method. Table 2 demonstrates that the three methods exhibit very close point and interval forecast accuracies.

Table 2: Comparison of point and interval forecasting accuracies of TS methods with standard functional principal component analysis, robust functional principal component analysis and robust methods utilising the regularised singular value decomposition, denoted as FPCA, M-FPCA and RobRSVD, respectively.

	MSFE			MAFE			Interval score		
	FPCA	M-FPCA	RobRSVD	FPCA	M-FPCA	RobRSVD	FPCA	M-FPCA	RobRSVD
Min	0.7939	0.7768	0.7753	0.6900	0.6774	0.6770	8.2059	8.2445	8.6548
Median	8.1868	8.0633	8.1196	2.3820	2.3789	2.3792	9.1384	9.0990	9.2289
Mean	22.0173	22.0134	22.0224	3.0447	3.0429	3.0465	15.7863	15.8190	15.6758
Max	354.7139	355.0930	354.9265	16.1869	16.1776	16.1833	128.4479	130.4321	127.0060

7.1 Updating point forecasts

As we receive partially observed data from the most recent curve, we can dynamically update our forecasts in the hope of achieving better forecast accuracy. Using the five dynamic updating methods detailed in Section 4, we construct plots of averaged MSFE and MAFE across all discretised (evaluation) time points as shown in Figure 3. When the BM method, the RR method, and functional linear regression are used, the forecast errors tend to decrease with some fluctuation. Averaged over the forecasting period and measured by the averaged MAFE, the functional linear regression has the best point forecast accuracy. Assessed by the averaged MSFE, the functional linear regression also has the lowest forecast errors across all daily trading times, except for the hour (approx.) from 13:00:00 to 15:00:00 during which time the most volatile movements of the VIX index were recorded as shown in Figure 3. As we receive increasingly more partially observed data in the most recent curve, dynamic updating methods tend to perform much better than the TS method. These unsurprising results highlight the advantages of a functional data-analytic approach that views data as realisations of a continuous stochastic process. As the frequency of the data increases, the performance of the dynamic updating method should improve even further, as interpolation errors will likely decrease.

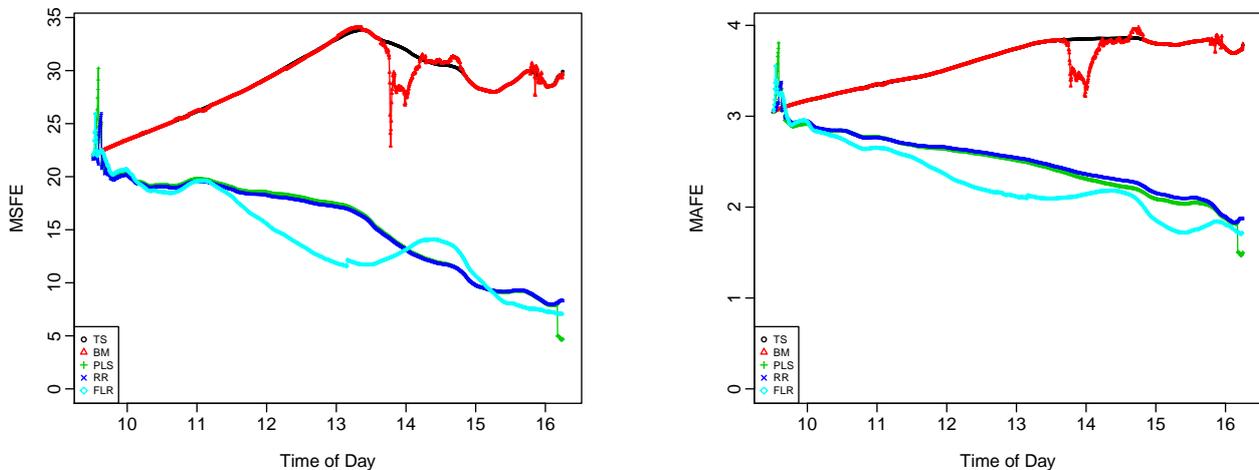


Figure 3: A comparison of point forecast accuracy between the TS method and four dynamic updating methods, as measured by MSFE and MAFE. Note that FLR stands for functional linear regression.

Diebold and Mariano (2002) tests are also conducted to formally assess the most accurate forecasting method, with hypotheses and corresponding results shown in Table 3. For example, in Test 1 we consider a squared-loss function and an absolute-loss function as defined in Equation (16). The null hypothesis is that the two forecasts have the same accuracy, while the alternative hypothesis is that the FLR method is more accurate than the TS model. With the help of the *forecast* package in R (R Core Team 2018), almost zero p -values ($< 2.2e - 16$) are obtained for both loss functions, indicating that the null hypothesis should be rejected. Results in all other

rows of the same table can be interpreted similarly. Overall, Table 3 shows that the functional linear regression model outperforms all other methods in producing accurate point forecasts. Such superior performance is not only sample-specific but can also draw inferences on the entire population.

Table 3: Diebold-Mariano tests of accuracy for point forecasts obtained by the TS method and four dynamic updating methods.

	Hypotheses	Squared-error loss		Absolute-error loss	
		Test Statistic	p -value	Test Statistic	p -value
Test 1	$H_0 : \text{Error}_{\text{FLR}} = \text{Error}_{\text{TS}}$ $H_A : \text{Error}_{\text{FLR}} < \text{Error}_{\text{TS}}$	-89.0106	$< 2.2e - 16$	-83.9062	$< 2.2e - 16$
Test 2	$H_0 : \text{Error}_{\text{FLR}} = \text{Error}_{\text{BM}}$ $H_A : \text{Error}_{\text{FLR}} < \text{Error}_{\text{BM}}$	-88.7233	$< 2.2e - 16$	-83.3199	$< 2.2e - 16$
Test 3	$H_0 : \text{Error}_{\text{FLR}} = \text{Error}_{\text{PLS}}$ $H_A : \text{Error}_{\text{FLR}} < \text{Error}_{\text{PLS}}$	-24.1949	$< 2.2e - 16$	-43.4326	$< 2.2e - 16$
Test 4	$H_0 : \text{Error}_{\text{FLR}} = \text{Error}_{\text{RR}}$ $H_A : \text{Error}_{\text{FLR}} < \text{Error}_{\text{RR}}$	-21.7857	$< 2.2e - 16$	-49.3449	$< 2.2e - 16$

Forecast results reported so far are obtained for the VIX index data from January to June of 2017. To add robustness to the results, we have also conducted an extended analysis using all of 2016 (details included in the Supplement), providing additional support for the proposition that functional time series forecasting models with dynamic updating (especially the FLR model) should be considered in order to produce accurate forecasts of the VIX index. Furthermore, it also indicates that our proposed models do not depend on a particular data range to provide the desired results outlined here.

7.2 Updating interval forecasts

After observing the VIX intraday cumulative returns from 9:30:00 to 14:59:45 on 30 June 2017, the BM, PLS, and the functional linear regression method can be applied to dynamically update the interval forecasts for the remaining trading time period of the same day. Figure 4 shows the 80% pointwise prediction intervals for the next hour using three dynamic updating methods, compared with the TS model that does not incorporate updating. It can be seen that the functional linear regression produces the narrowest prediction interval, leading to the most accurate evaluation of forecast uncertainty.

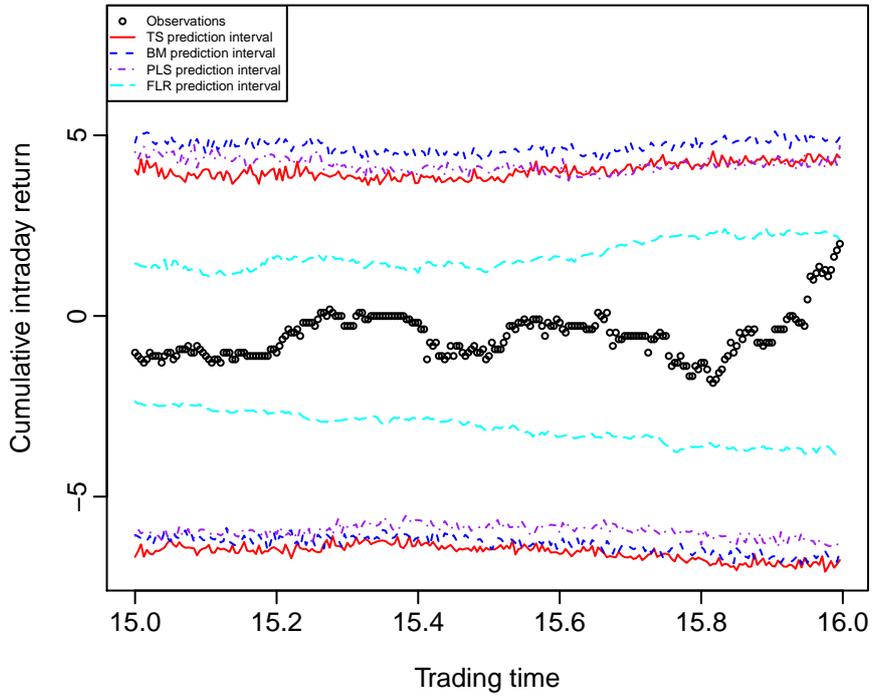


Figure 4: Updating 80% pointwise prediction intervals using the TS method, BM method, PLS method, and functional linear regression, respectively.

7.3 Economic evaluation of forecasts

Building on the superior predictability of dynamic functional time series models, we consider applications that highlight the economic value of adopting the proposed forecasting methods. As the VIX index itself cannot be directly traded, implementing a trading strategy involves the adoption of intraday data for a related derivative product. To the best of our knowledge, there are no existing intraday VIX trading strategies outlined in the literature,⁶

Utilising the MME lose functions and MCPDC measure, an indication of how well our VIX forecasts at providing trading signals is presented in Figure 5. Our proposed functional time series methods all have MCPDC⁷ values that are better than the commonly used benchmark. By benchmark, we mean that a random model should correctly predict the direction of change of the index 50% of the time. Among our functional time series methods, the FLR model produces the highest mean MCPDC of about 80%. This result indicates that our considered functional time series approaches should perform well when adapted to provide a buy/sell indicator to pursue long/short positions in VIX index derivative products. Furthermore, the Mean Mixed Error (Under/Over) results indicate that in general there is no significant bias arising from our approaches, i.e., the forecasts do not systematically under or over predict the VIX index level. Such

⁶For instance, Konstantinidi et al. (2008), Konstantinidi and Skiadopoulos (2011) and Kourtis et al. (2016) only consider trading strategies that were utilising either daily or monthly futures data.

⁷MCPDC of TS: max 63.75%, mean 60.48%; MCPDC of BM: max 70.47%, mean 60.88%; MCPDC of PLS: max 88.50%, mean 75.22%; MCPDC of RR: max 85.75% mean 76.98%; MCPDC of FLR: max 88.95% mean 79.13%.

a concern would be particularly relevant for practitioners with stop-loss limits, or those informing the purchase of VIX-related options using our VIX index forecasts.

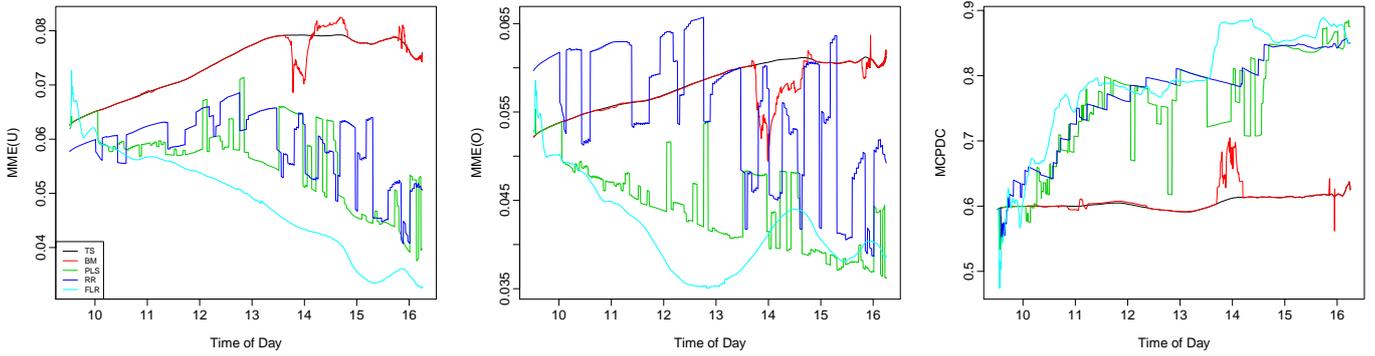


Figure 5: Mean Mixed Error (Under), labelled as $MME(U)$; Mean Mixed Error (Over), labelled as $MME(O)$; and Mean Correct Predictor of the Direction of Change, labelled as $MCPDC$, computed for the TS method and four dynamic updating methods.

However, many practical considerations need to be accounted for when implementing any trading strategy. These factors are particularly numerous in a high-frequency environment, such as ours. Many are idiosyncratic to individual market participants including, market frictions, slippage, transaction costs, order book depth, microstructure noise, order latency, exchange rebates, stop-loss limits, and leverage and portfolio diversification concerns. Without all knowledge of such factors, outlining a simulated trading strategy would prove to be unrealistic.

8 Conclusion

Functional time series forecasting and updating methods discussed here treat observed data as realisations of a continuous stochastic process. We use functional PCA to model the temporal dependence among functional curves, reducing the dimensionality of data. A set of functional principal components that explains at least 90% of the total variation in the 15-second VIX data is considered so that the main features of the original functional time series can be maintained. We produce one-day-ahead forecast VIX CIDRs by forecasting K retained principal component scores through a univariate or multivariate time series forecasting technique. In particular, the forecasts are obtained by multiplying the forecast principle component scores by the estimated functional principal components and then adding the estimated mean function.

We consider five dynamic updating methods, which utilise partially observed data for the most recent curve to improve forecast accuracy. The BM method applies the general structure of the TS method to our VIX data but rearranges the function support to update forecasts during the trading day. The OLS method treats the partially observed data in the most recent curve as a discrete response vector and regresses it against the corresponding discretised functional

principal components. However, the singularity problem may plague the implementation of the OLS method, when the number of partially observed data points is less than the number of functional principal components. To overcome this problem, we consider methods that shrink the estimated regression coefficients. The RR method shrinks the regression coefficient estimates towards zero, whereas the PLS method shrinks the regression coefficient estimates towards $\hat{\beta}_{n+1|n}^{\text{TS}}$. The amount of shrinkage in the RR and PLS methods is controlled by a shrinkage parameter λ that can be tuned by minimising the forecast error measure within a holdout data sample. One commonality between the OLS, RR, and PLS methods is that they attempt to update forecasts by utilising discrete time points. Conversely, the functional linear regression first decomposes two blocks of functional time series that correspond to the partially observed data and remaining data periods via the Karhunen-Loève expansion. The method then forms a linear model based on the two sets of principal component scores via the ordinary least squares method and estimates the regression coefficient function from which the updated forecasts can be obtained. Functional linear regression outperforms all other dynamic updating methods investigated in producing the most accurate point forecasts assessed using averaged MSFE and MAFE over the different discretised time points in the testing sample.

As a means of assessing forecast uncertainty, a non-parametric bootstrap method is used to construct prediction intervals for the TS and BM methods. The PLS prediction intervals are updated as we sequentially observe new data in the most recent curve. When updating prediction intervals through functional linear regression, we considered the maximum entropy bootstrapping and the non-parametric bootstrapping method to obtain bootstrapped regression coefficient functions and bootstrapped error functions, respectively. These approaches result in improved interval forecast accuracy, allowing us to compare the different functional time series methods considered.

The focus of this paper is on statistical forecasting accuracy. However, the real-time dynamically updating VIX forecasting results bring with them a number of important economic implications. Firstly, as shown by [Psaradellis and Sermpinis \(2016\)](#), implied volatility index forecasts can feed into profitable trading strategies involving tradable VIX derivatives, which despite many additional considerations, could also be adopted by practitioners in our high frequency environment. Secondly, considering prediction intervals as we do, boasts the advantage of enabling a number of different strategies to be planned and specified for a range of possible scenarios ([Chatfield 1993](#)). Lastly, demonstrating predictability in the ‘fear gauge’ is of great benefit to market analysts and practitioners alike, as it enables them to better prepare for the impact of future market movements. For these reasons, economic evaluations of hedging and trading strategies based on the proposed functional time series methods could form the basis for future research.

Acknowledgements

The authors thank insightful comments and suggestions from three reviewers. The first author acknowledges a faculty research grant from the College of Business and Economics at the Australian National University. The second author would like to acknowledge financial support from the Australian Government Research Training Program Stipend Scholarship.

References

- Aneiros-Pérez, G. and Vieu, P. (2008), 'Nonparametric time series prediction: A semi-functional partial linear modeling', *Journal of Multivariate Analysis* **99**(5), 834–857.
- Aue, A., Norinho, D. D. and Hörmann, S. (2015), 'On the prediction of stationary functional time series', *Journal of the American Statistical Association* **110**(509), 378–392.
- Bali, T. G. and Theodossiou, P. (2007), 'A conditional-SGT-VaR approach with alternative GARCH models', *Annals of Operations Research* **151**(1), 241–267.
- Bernales, A. and Guidolin, M. (2014), 'Can we forecast the implied volatility surface dynamics of equity options? predictability and economic value tests', *Journal of Banking & Finance* **46**, 326–342.
- Besse, P., Cardot, H. and Stephenson, D. (2000), 'Autoregressive forecasting of some functional climatic variations', *Scandinavian Journal of Statistics* **27**(4), 673–687.
- Bosq, D. (2000), *Linear Processes in Function Spaces*, Lecture notes in Statistics, New York.
- Brailsford, T. J. and Faff, R. W. (1996), 'An evaluation of volatility forecasting techniques', *Journal of Banking & Finance* **20**(3), 419–438.
- Chatfield, C. (1993), 'Calculating interval forecasts', *Journal of Business & Economic Statistics* **11**(2), 121–135.
- Chiou, J.-M. et al. (2012), 'Dynamical functional prediction and classification, with application to traffic flow prediction', *The Annals of Applied Statistics* **6**(4), 1588–1614.
- Diebold, F. X. and Mariano, R. S. (2002), 'Comparing predictive accuracy', *Journal of Business & economic statistics* **20**(1), 134–144.
- Dunis, C. L., Likothanassis, S. D., Karathanasopoulos, A. S., Sermpinis, G. S. and Theofilatos, K. A. (2013), 'A hybrid genetic algorithm–support vector machine approach in the task of forecasting and trading', *Journal of Asset Management* **14**(1), 52–71.
- Febrero, M., Galeano, P. and González-Manteiga, W. (2007), 'A functional analysis of NOx levels: Location and scale estimation and outlier detection', *Computational Statistics* **22**(3), 411–427.
- Fernandes, M., Medeiros, M. C. and Scharth, M. (2014), 'Modeling and predicting the CBOE market volatility index', *Journal of Banking & Finance* **40**, 1–10.
- Fuertes, A.-M., Izzeldin, M. and Kalotychou, E. (2009), 'On forecasting daily stock volatility: The role of intraday information and market conditions', *International Journal of Forecasting* **25**(2), 259–281.

- Gabriel, K. R. and Zamir, S. (1979), 'Lower rank approximation of matrices by least squares with any choice of weights', *Technometrics* **21**(4), 489–498.
- Gabrys, R., Horváth, L. and Kokoszka, P. (2010), 'Tests for error correlation in the functional linear model', *Journal of the American Statistical Association* **105**(491), 1113–1125.
- Gneiting, T. and Katzfuss, M. (2014), 'Probabilistic forecasting', *Annual Review of Statistics and Its Application* **1**, 125–151.
- Gneiting, T. and Raftery, A. E. (2007), 'Strictly proper scoring rules, prediction, and estimation', *Journal of the American Statistical Association* **102**(477), 359–378.
- Hall, P. (2011), Principal component analysis for functional data: Methodology, theory and discussion, in F. Ferraty and Y. Romain, eds, 'The Oxford Handbook of Functional Data Analysis', Oxford: Oxford University Press, pp. 210–234.
- Hall, P. and Vial, C. (2006), 'Assessing the finite dimensionality of functional data', *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **68**(4), 689–705.
- Hays, S., Shen, H. and Huang, J. Z. (2012), 'Functional dynamic factor models with application to yield curve forecasting', *The Annals of Applied Statistics* **6**(3), 870–894.
- Hoerl, A. E. and Kennard, R. W. (1970), 'Ridge regression: Biased estimation for nonorthogonal problems', *Technometrics* **12**(1), 55–67.
- Hörmann, S., Horváth, L. and Reeder, R. (2013), 'A functional version of the ARCH model', *Econometric Theory* **29**(2), 267–288.
- Horváth, L. and Kokoszka, P. (2012), *Inference for Functional Data with Applications*, Springer, New York.
- Huang, J. Z., Shen, H. and Buja, A. (2009), 'The analysis of two-way functional data using two-way regularized singular value decompositions', *Journal of the American Statistical Association* **104**(488), 1609–1620.
- Huber, P. J. and Ronchetti, E. M. (2009), *Robust Statistics*, John Wiley & Sons, Hoboken, New Jersey.
- Hubert, M., Rousseeuw, P. J. and Verboven, S. (2002), 'A fast method for robust principal components with applications to chemometrics', *Chemometrics and Intelligent Laboratory Systems* **60**(1), 101–111.
- Hyndman, R., Athanasopoulos, G., Bergmeir, C., Caceres, G., Chhay, L., O'Hara-Wild, M., Petropoulos, F., Razbash, S., Wang, E. and Yasmien, F. (2018), *forecast: Forecasting functions for time*

series and linear models. R package version 8.4.

URL: <https://CRAN.R-project.org/package=forecast>

- Hyndman, R. J. and Khandakar, Y. (2008), 'Automatic time series for forecasting: the forecast package for R', *Journal of Statistical Software* **27**(3).
- Hyndman, R. J. and Shang, H. L. (2009), 'Forecasting functional time series (with discussions)', *Journal of the Korean Statistical Society* **38**(3), 199–221.
- Hyndman, R. J. and Shang, H. L. (2010), 'Rainbow plots, bagplots, and boxplots for functional data', *Journal of Computational and Graphical Statistics* **19**(1), 29–45.
- Hyndman, R. J. and Ullah, M. S. (2007), 'Robust forecasting of mortality and fertility rates: A functional data approach', *Computational Statistics & Data Analysis* **51**(10), 4942–4956.
- Karhunen, K. (1946), 'Zur spektraltheorie stochastischer prozesse', *Ann. Acad. Sci. Fennicae, AI* **34**.
- Kearney, F., Cummins, M. and Murphy, F. (2018), 'Forecasting implied volatility in foreign exchange markets: a functional time series approach', *The European Journal of Finance* **24**(1), 1–18.
- Kılıç, D. K. and Uğur, O. (2018), 'Multiresolution analysis of S&P500 time series', *Annals of Operations Research* **260**(1-2), 197–216.
- Klepsch, J. and Klüppelberg, C. (2016), An innovations algorithm for the prediction of functional linear processes, Working paper, Technische Universität München.
URL: <https://arxiv.org/abs/1607.05874>
- Klepsch, J., Klüppelberg, C. and Wei, T. (2017), 'Prediction of functional ARMA processes with an application to traffic data', *Econometrics and Statistics* **1**, 128–149.
- Kokoszka, P., Miao, H. and Zhang, X. (2014), 'Functional dynamic factor model for intraday price curves', *Journal of Financial Econometrics* **13**(2), 456–477.
- Kokoszka, P. and Reimherr, M. (2013), 'Determining the order of the functional autoregressive model', *Journal of Time Series Analysis* **34**(1), 116–129.
- Kokoszka, P. and Reimherr, M. (2017), *Introduction to Functional Data Analysis*, CRC Press, Boca Raton, FL.
- Kokoszka, P., Rice, G. and Shang, H. L. (2017), 'Inference for the autocovariance of a functional time series under conditional heteroscedasticity', *Journal of Multivariate Analysis* **162**, 32–50.
- Konstantinidi, E. and Skiadopoulos, G. (2011), 'Are VIX futures prices predictable? An empirical investigation', *International Journal of Forecasting* **27**(2), 543–560.

- Konstantinidi, E., Skiadopoulos, G. and Tzagkaraki, E. (2008), 'Can the evolution of implied volatility be forecasted? Evidence from European and US implied volatility indices', *Journal of Banking & Finance* **32**(11), 2401–2411.
- Kourtis, A., Markellos, R. N. and Symeonidis, L. (2016), 'An international comparison of implied, realized, and GARCH volatility forecasts', *Journal of Futures Markets* **36**(12), 1164–1193.
- Li, D., Robinson, P. M. and Shang, H. L. (2018), 'Long-range dependent curve time series', *Journal of the American Statistical Association: Theory and Methods* **under revision**.
- Loève, M. (1946), 'Fonctions aléatoires à décomposition orthogonale exponentielle', *La Revue Scientifique* **84**, 159–162.
- Maronna, R., Martin, R. D. and Yohai, V. (2006), *Robust Statistics*, John Wiley & Sons, Chichester, West Sussex, England.
- Morris, J. S. (2015), 'Functional regression', *Annual Review of Statistics and Its Application* **2**, 321–359.
- Müller, H.-G. (2005), 'Functional modelling and classification of longitudinal data', *Scandinavian Journal of Statistics* **32**(2), 223–240.
- Müller, H.-G., Sen, R. and Stadtmüller, U. (2011), 'Functional data analysis for volatility', *Journal of Econometrics* **165**(2), 233–245.
- Peña, D. and Sánchez, I. (2007), 'Measuring the advantages of multivariate vs. univariate forecasts', *Journal of Time Series Analysis* **28**(6), 886–909.
- Psaradellis, I. and Sermpinis, G. (2016), 'Modelling and trading the U.S. implied volatility indices. Evidence from the VIX, VXN and VXD indices', *International Journal of Forecasting* **32**(4), 1268 – 1283.
- R Core Team (2018), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria.
URL: <https://www.R-project.org/>
- Ramsay, J. O. and Silverman, B. W. (2005), *Functional Data Analysis*, 2nd edn, Springer, New York.
- Rice, J. A. and Silverman, B. W. (1991), 'Estimating the mean and covariance structure nonparametrically when the data are curves', *Journal of the Royal Statistical Society. Series B (Methodological)* **1**, 233–243.
- Sermpinis, G., Stasinakis, C., Theofilatos, K. and Karathanasopoulos, A. (2014), 'Inflation and unemployment forecasting with genetic support vector regression', *Journal of Forecasting* **33**(6), 471–487.

- Sguera, C., Galeano, P. and Lillo, R. E. (2016), 'Functional outlier detection by a local depth with application to NO_x levels', *Stochastic environmental research and risk assessment* **30**(4), 1115–1130.
- Shang, H. L. (2014), 'A survey of functional principal component analysis', *AStA Advances in Statistical Analysis* **98**(2), 121–142.
- 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. (2018), 'Bootstrap methods for stationary functional time series', *Statistics and Computing* **28**(1), 1–10.
- Shang, H. L. and Hyndman, R. J. (2011), 'Nonparametric time series forecasting with dynamic updating', *Mathematics and Computers in Simulation* **81**(7), 1310–1324.
- Shen, H. (2009), 'On modeling and forecasting time series of smooth curves', *Technometrics* **51**(3), 227–238.
- Tsay, R. S. (2013), *Multivariate Time Series Analysis: with R and Financial Applications*, John Wiley & Sons, Hoboken, New Jersey.
- Vinod, H. D. (2004), 'Ranking mutual funds using unconventional utility theory and stochastic dominance', *Journal of Empirical Finance* **11**(3), 353–377.
- Vinod, H. D., López-de Lacalle, J. et al. (2009), 'Maximum entropy bootstrap for time series: the meboot R package', *Journal of Statistical Software* **29**(5), 1–19.
- Wang, L. and Zhu, J. (2010), 'Financial market forecasting using a two-step kernel learning method for the support vector regression', *Annals of Operations Research* **174**(1), 103–120.
- Whaley, R. E. (2000), 'The investor fear gauge', *The Journal of Portfolio Management* **26**(3), 12–17.
- Xu, J. (1999), 'Modeling Shanghai stock market volatility', *Annals of Operations Research* **87**(0), 141–152.
- Yao, F., Müller, H.-G. and Wang, J.-L. (2005), 'Functional data analysis for sparse longitudinal data', *Journal of the American Statistical Association* **100**(470), 577–590.
- Zhang, L., Shen, H. and Huang, J. Z. (2013), 'Robust regularized singular value decomposition with application to mortality data', *The Annals of Applied Statistics* **7**(3), 1540–1561.