AMSI VACATION RESEARCH SCHOLARSHIPS 2022–23

SET YOUR SIGHTS ON RESEARCH THIS SUMMER

Predicting Stock Price Volatility within a Month – A Functional Data Analysis Approach

Runqiu Fei

Supervised by Guoqi Qian

The University of Melbourne

Abstract

Modelling and predicting stock price volatility is one of the key research fields for risk management in today's financial stock market. In support of today's frequent and continuous decision makings in the market, more sophisticated and informative forecasting techniques are strongly desired. In this project, we employ Functional Data Analysis (FDA) to study intra-month stock price volatility. FDA considers covariates as functional curves and gives response forecasting as functional trajectories inheriting valuable functional natures including flexibility and information richness. In particular, for our project, we use a functional concurrent AR(1) model to predict the intra-month volatility trajectory of this month from that of last month. We introduce some essential FDA techniques including curve smoothing, functional model fitting and functional F-test in the context of volatility forecasting. Then, we empirically apply FDA to a Dow Jones Industrial Average (DJIA) dataset to do intra-month volatility forecasting. We find that FDA is empirically effective in intra-month volatility forecasting. Further, the forecasted functional trajectories should be highly informative and hence preferable in today's fast-moving markets.

1 Introduction

The financial stock market has always been an attractive place for Mathematical and Statistical research due to its stochastic nature and richness of data. Typically, modelling and predicting stock price volatility is one of the key research fields for risk management. Understanding the nature of volatility contributes to the development of economic theories (e.g., see Binder and Merges (2001)). Meanwhile, accurate volatility prediction helps policymakers to better design the market (e.g., see Ivrendi and Guloglu (2012)) and traders to invest wisely (e.g., see Huber, Huber and Kirchler (2022)).

Studying volatility is indeed rewarding, but also challenging. One problem is that the trading days are not evenly distributed over a time horizon, which makes it hard to model daily volatilities using traditional models for evenly spaced time series (e.g., see the GARCH model in Bollerslev (1986)). Another problem is how to fully utilise the data to support frequent and continuous decision makings (e.g., see High-Frequency Trading in Conerly (2014)). Given such high velocity and large volume of today's stock market data, people are no longer satisfied with a single mean prediction. Instead, a continuous path or trajectory prediction should be much more informative and hence preferable.

Targeting these two challenges, our project aims to employ Functional Data Analysis (FDA) proposed by Ramsay and Silverman (2005) to study and predict the intra-month trajectory of stock price volatility. Specifically, inspired by Alva, Romo, and Ruiz Ortega (2009), we adapt a functional concurrent AR(1) model to predict the intra-month volatility trajectory of the next month from that of the current month. Note that to be consistent with Alva, Romo, and Ruiz Ortega (2009), a daily volatility observation is defined as the absolute logarithmic returns of the day. The key idea of FDA in the context is curve smoothing. By smoothing the daily volatility observations within a month into a continuous functional curve, the uneven time spaces between trading days become irrelevant. Moreover, by considering the response as a smooth functional curve, much more information will be extractable including the volatility at any given time point within the domain or even some meaningful derivatives of the curve.

The rest of the report is structured as follows. In section 2, we introduce essential FDA techniques and the functional concurrent AR(1) model in the context of intra-month volatility prediction. In section 3, we apply the functional concurrent AR(1) model to the Dow Jones Industrial Average (DJIA) data to empirically examine the effectiveness of the model. Finally, we conclude and discuss our research findings in section 4.

Statement of Authorship

Conceptualisation, Guoqi Qian (G.Q.) and Runqiu Fei (R.F.); methodology, G.Q. and R.F.; software, R.F.; validation, G.Q. and R.F.; formal analysis, G.Q. and R.F.; investigation, G.Q. and R.F.; writing—original draft preparation, R.F.; writing—review and editing, G.Q. and R.F.; visualisation, R.F.; supervision, G.Q.

2 Method

In this section, we first rigorously define stock price volatility. Then, we introduced some basic curve smoothing techniques and functional concurrent AR(1) model in the context of volatility prediction. In addition, some prediction and validation methods will be discussed.

2.1 Stock Price Volatility

We would define stock price volatility similar to that in Alva, Romo, and Ruiz Ortega (2009) for consistency. In our research, volatilities are observed daily for each trading day within each month, and it is defined as the absolute logarithmic returns of the day. More rigorously, for trading day j of month i, the daily volatility is define as

$$v_{i,j} = \left| \log \frac{p_{i,j}}{p_{i,j-1}} \right|, \quad i = 1, ..., n, \quad j = 2, ..., T_i, \tag{1}$$

where $p_{i,j}$ denotes the (average) price of the stock in trading day j of month i. Note that n is the number of months in our sample and T_i is the number of trading days in month i. Also note that volatilities are observed from trading day 2 onwards for each month, so there will be $(T_i - 1)$ volatility observations for month i. Intuitively, a daily volatility $v_{i,j} = 0$ only when the price stays unchanged (i.e., $p_{i,j} = p_{i,j-1}$); otherwise, $v_{i,j} > 0$ and will increase as the price (absolute) difference increases. Then, an intra-month volatility functional curve $V_i(t)$ for month i can be obtained by smoothing the observations $v_{i,j}$, $j = 2, ..., T_i$ within month i into a continuous and smooth curve. In this case, $V_1(t)$,...,$V_n(t)$ can be regarded as n samples from the intra-month volatility function $V(t)$, and we are building functional AR(1) model based on the samples.

Note that since the number of trading days T_i may vary from month to month, the domain unification issue should be properly dealt with. In our project, we unify the domains of $V_i(t)$, $i = 1, ..., n$ to a common domain $[0, 1]$ by mapping trading day j of month i to time $t_{i,j} = \frac{j-2}{T_i-2} \in [0, 1]$, $j = 2, ..., T_i$. That is, we evenly distribute the trading days of a month across the interval $[0, 1]$, with the first day mapping to 0 and the last day mapping to 1. In this case, it is appropriate to assume $[0, 1]$ as the domain of the intra-month volatility function $V(t)$. This domain unification procedure is actually an example of landmark registration (Ramsay, Hooker and Graves 2009), which is a commonly used FDA technique for curve alignment.

2.2 Curve Smoothing

We would take the first-month volatility function $V_1(t)$ as an example to illustrate how curve smoothing in FDA works in this context. In FDA, we appriximate $V_1(t)$ by a linear combination of a functional basis system $\boldsymbol{\phi}(t) = (\phi_1(t), ..., \phi_K(t))^T$. That is,

$$V_1(t) = \sum_{k=1}^{K} c_k \phi_k(t) + \eta_1(t) = \mathbf{c}^T \boldsymbol{\phi}(t) + \eta_1(t), \tag{2}$$

where $\mathbf{c} = (c_1, ..., c_K)^T$ are the coefficients to be determined, K is the number of basis functions in the basis system and $\eta_1(t) \sim Normal(0, \sigma^2(t))$ is a random error. In our project, we use a cubic B-spline basis system, which is commonly used in general settings (Ramsay, Hooker and Graves 2009). The next question is how to choose the number of basis functions K. Normally, $K << (T_1 - 1)$ is chosen to prevent potential overfitting (Ramsay, Hooker and Graves 2009). But underfitting can also be a problem if K is too small. Actually, choosing an optimal K is not fundamental. Instead, it is a common practice to first choose a relatively large K, and then apply a roughness penalty to adjust the potential overfitting.

That is, we perform curve smoothing with roughness penalties (Ramsay, Hooker and Graves 2009). Specifically, we find the best set of coefficients $\mathbf{c}$ that minimises the smoothing criterion function

$$F(\mathbf{c}) = \text{SSE}(V_1) + \lambda \text{PEN}_2(V_1) = \sum_{j=2}^{T_1} [v_{1,j} - V_1(t_j)]^2 + \lambda \int_0^1 [D^2 V_1(t)]^2 \, dt, \tag{3}$$

where a smaller sum of squared errors $\text{SSE}(V_1)$ indicates a better fit, and $\text{PEN}_2(V_1)$ is a penalty term penalising the curvature (i.e., second derivative) of $V_1(t)$ to prevent overfitting. The smoothing parameter $\lambda > 0$ can be chosen by minimising a generalized cross-validation measure $\text{GCV}(\lambda)$ proposed by Craven and Wahba (1979), but we would not go into details here since it is not the primary focus of our project.

In R, we can use the function `create.bspline.basis` from the `fda` package to create the required B-spline basis system, and then use the functions `fdPar` and `smooth.basis` from the same package to estimate the coefficients $\mathbf{c}$ and obtain the smoothed functional curves. See more details in Ramsay, Hooker and Graves (2009).

2.3 Functional Concurrent AR(1) Model

Inspired by Alva, Romo, and Ruiz Ortega (2009) who employed FDA for intra-day volatility prediction, we propose the following functional concurrent AR(1) model to predict the intra-month volatility function (trajectory) of month i from that of month $(i - 1)$

$$V_i(t) = \beta_0(t) + \beta_1(t) V_{i-1}(t) + \epsilon_i(t), \tag{4}$$

where $\epsilon_i(t)$ is the error term. This model is called *concurrent* (Ramsay, Hooker and Graves 2009) since it only relates the value of $V_i(t)$ to the value of $V_{i-1}(t)$ at the same time points t. Actually, this concurrent AR(1) model may not be the finest model for $V_i(t)$. The volatility $V_i(t)$ at a given time point t of the current month may also correlate with $V_{i-1}(s)$ where $s \neq t$ (i.e., volatilities at other time points s of last month), or with $V_i(h)$ where $h < t$ (i.e., volatilities at earlier time points h of the current month). Other factors such as seasonality may also have effects on $V_i(t)$. Nevertheless, the concurrent AR(1) model has its own advantage that it is simpler to understand and intuitively easier to interpret, which is why we consider it preferable for this time-limited project.

The key process in model fitting is parameter estimation. For this functional concurrent AR(1) model, the estimation procedure is largely similar to that of general concurrent models (Ramsay, Hooker and Graves 2009). We first try to express the model in matrix notation. Suppose there are N pairs of the functional observations $(V_i(t), V_{i-1}(t))$, $i = 2, ..., N+1$. Let the $N \times 2$ functional matrix $\mathbf{Z}(t)$ contain all 1's in the first column and the covariate functions $V_{i-1}(t)$ in the second column. And let the 2×1 functional coefficient vector $\boldsymbol{\beta}(t) = (\beta_0(t), \beta_1(t))^T$. Then, the functional concurrent AR(1) model in matrix notation is

$$\mathbf{V}(t) = \mathbf{Z}(t)\boldsymbol{\beta}(t) + \boldsymbol{\epsilon}(t), \tag{5}$$

where $\mathbf{V}(t)$ is an $N \times 1$ functional vector containing the response functions $V_i(t)$. Now, we can express the corresponding $N \times 1$ vector of residual functions as

$$\mathbf{r}(t) = \mathbf{V}(t) - \mathbf{Z}(t)\boldsymbol{\beta}(t). \tag{6}$$

Finally, our target is to find the best functional coefficient vector $\boldsymbol{\beta}(t)$ that minimises the weighted regularised fitting criterion

$$\text{LMSSE}(\boldsymbol{\beta}) = \int_0^1 \mathbf{r}(t)^T \mathbf{r}(t)\, dt + \sum_{j=0}^{1} \lambda_j \int_0^1 [D^2 \beta_j(t)]^2\, dt, \tag{7}$$

where the first term is a functional equivalent SSE measuring the goodness of fit and the second term is a penalty term penalising the curvatures (i.e., second derivatives) of $\beta_j(t)$, $j = 0, 1$ to prevent overfitting. Note that we would also approximate $\beta_j(t)$, $j = 0, 1$ as linear combinations of functional bases and perform similar curve smoothing procedures.

In R, the function `fRegress` from the `fda` package can be used to fit a functional concurrent linear model (including our concurrent AR(1) model), and estimate the relevant functional coefficients $\boldsymbol{\beta}(t)$. See more details in Ramsay, Hooker and Graves (2009).

2.4 Functional Permutation F-Test

We would use a functional permutation F-test proposed by Ramsay, Hooker and Graves (2009) to examine the significance of the predictive relationship between this-month volatility (i.e., the covariate) and next-month volatility (i.e., the response). The idea of this permutation F-test is that if there is no relationship between the response and the covariate, then the fitted model would not vary significantly when we try to randomly pair up the response values with the covariate values. In other words, if the observed pairing of response and covariate values trains the model significantly better than other random pairings, then we say that the observed data support a significant predictive relationship between the covariate and the response. The goodness of fit of a model is measured by the following (functional) F-statistic

$$F(t) = \frac{\text{Var}[\hat{\mathbf{V}}(t)]}{\frac{1}{n}\sum_{i=1}^{n}(V_i(t) - \hat{V}_i(t))^2}, \tag{8}$$

where $\hat{\mathbf{V}}(t) = \{\hat{V}_i(t),\ i = 1, ..., n\}$ are the predicted functional response curves and a large F-statistic indicates a better fit. To form a (functional) null distribution, we iteratively calculate the (functional) F-statistics $F_p(t),\ p = 1, ..., m$, for m different models fitted respectively by m different (random) response-covariate pairings. Then, we would obtain a functional point-wise critical value curve $c(t)$ as the 95% (functional) quantile curve of the (functional) null distribution $\{F_p(t),\ p = 1, ..., m\}$. Then, the predictive relationship between the covariate and the response is statistically significant (at 5% level) if the (functional) F-statistic for the observed response-covariate pairing is point-wisely larger than the critical value curve (i.e., $F_{observed}(t) > c(t)$). More conservatively, we can obtain a scalar critical value c as the 95% (scalar) quantile of the (scalar) null distribution $\{\max_t F_p(t),\ p = 1, ..., m\}$. Then, if $F_{observed}(t) > c$, we conclude that the predictive relationship is statistically significant. Note that the scalar critical value c may be too conservative sometimes, in which case the functional critical value curve $c(t)$ should be more informative.

In R, the function `Fperm.fd` from the `fda` package can be used to perform a functional permutation F-test. See more details in Ramsay, Hooker and Graves (2009).

2.5 Volatility Forecasting and Validation

We would employ a multi-step forecasting (prediction) method to perform intra-month volatility forecasting. Suppose we have used smoothed volatility functions from n months (i.e., $\{V_i(t),\ i = 1, ..., n\}$) as a functional time series (with month n as the last month) to train our AR(1) model and we have obtained $\hat{\beta}_0(t), \hat{\beta}_1(t)$ as the functional coefficient estimates, then an m-step volatility forecasting works as follows:

1. Set $i := n + 1$, $\hat{V}_n(t) := V_n(t)$.
2. Forecast the volatility function of month i as $\hat{V}_i(t) := \hat{\beta}_0(t) + \hat{\beta}_1(t)\hat{V}_{i-1}(t)$.
3. If $i < m$, increment i by 1 and go to Step 2; else, finish forecasting.

After an m-step volatility forecasting, we would obtain m predicted volatility (functional) trajectories corresponding to m future months. That is, $\{\hat{V}_i(t),\ i = n + 1, ..., n + m\}$.

To further examine how effective our model is on intra-month volatility forecasting, we would compare the goodness of fit of the predicted functional curves to that of the smoothed functional curves. We would employ the weighted mean square error (WMSE) as a goodness of fit measure, which is widely used in many general settings. A smaller WMSE indicates a better fit for the observed data. Also note that we take the inverse distances as weights. That is, the weight for observed data point $v_{i,j}$ is $w_{i,j}^{predicted} = 1 \,/\, |v_{i,j} - \hat{V}_i(t_j)|$ for the predicted functional curve $\hat{V}_i(t)$ (or $w_{i,j}^{smoothed} = 1 \,/\, |v_{i,j} - V_i(t_j)|$ for the smoothed functional curve $V_i(t)$). This means that data points further from the functional curve have smaller weights, and hence we prefer the curve with a more robust fit.

3 Application

In this section, we apply the proposed functional concurrent AR(1) model to the Dow Jones Industrial Average (DJIA) dataset to empirically examine the effectiveness of the model on intra-month volatility trajectory forecasting. We perform the relevant curve smoothing, model fitting, statistical testing and forecasting validation, which will be illustrated along with intuitive visualisations.

3.1 DJIA Volatility Dataset

The Dow Jones Industrial Average (DJIA) is one of the most commonly used stock price indexes to represent the general stock price level in the U.S. stock market. In our project, we use a 5-year daily DJIA dataset, which can be retrieved online from Macrotrends (2023). Note that the daily DJIA data represents daily stock prices, so we need to add an additional step to obtain the corresponding daily volatilities based on the definition of volatility in section 2.1. Then, the obtained DJIA volatility dataset is the primary dataset for our experiment.

The DJIA volatility dataset consists of daily volatility data across 5 years (2018 - 2022). Since we are predicting intra-month volatility trajectories, we group the daily volatility observations by month so that the full dataset can be represented as $S = \{v_{i,j},\ i = 1, ..., 60,\ j = 2, ..., T_i\}$ where $v_{i,j}$ is the volatility of trading day j of month i as defined in section 2.1. Then, we use the subset from the first 4 years (2018 - 2021) as the training set (i.e., $S_{training} = \{v_{i,j},\ i = 1, ..., 48,\ j = 2, ..., T_i\}$), and the subset from the last 1 year (2022) as the test set(i.e., $S_{test} = \{v_{i,j},\ i = 49, ..., 60,\ j = 2, ..., T_i\}$). As a common practice, we use the training set to train the functional concurrent AR(1) model, and then use the test set to validate the effectiveness of the model on forecasting.

3.2 Curve Smoothing with DJIA Volatilities

We do curve smoothing for both the training set $S_{training}$ and the test set S_{test}, following the method specified in section 2.2. After the curve smoothing, we obtain the training set with smoothed functional observations (i.e., $S_{training}^{smoothed} = \{V_i(t),\ i = 1, ..., 48\}$ where $V_i(t)$ is the smoothed volatility function of month i), and similarly the smoothed test set $S_{test}^{smoothed} = \{V_i(t),\ i = 49, ..., 60\}$. Note that along with the curve smoothing, we unify the domains of $V_i(t),\ i = 1, ..., 60$ to a common domain $[0, 1]$, following the method discussed in section 2.1.

To have a more intuitive sense of the data and how curve smoothing works, we give a visualisation of the test set volatility data points $v_{i,j} \in S_{test}$, superposed by the corresponding smoothed intra-month volatility functional curves $V_i(t) \in S_{test}^{smoothed}$, as shown in Figure 1.

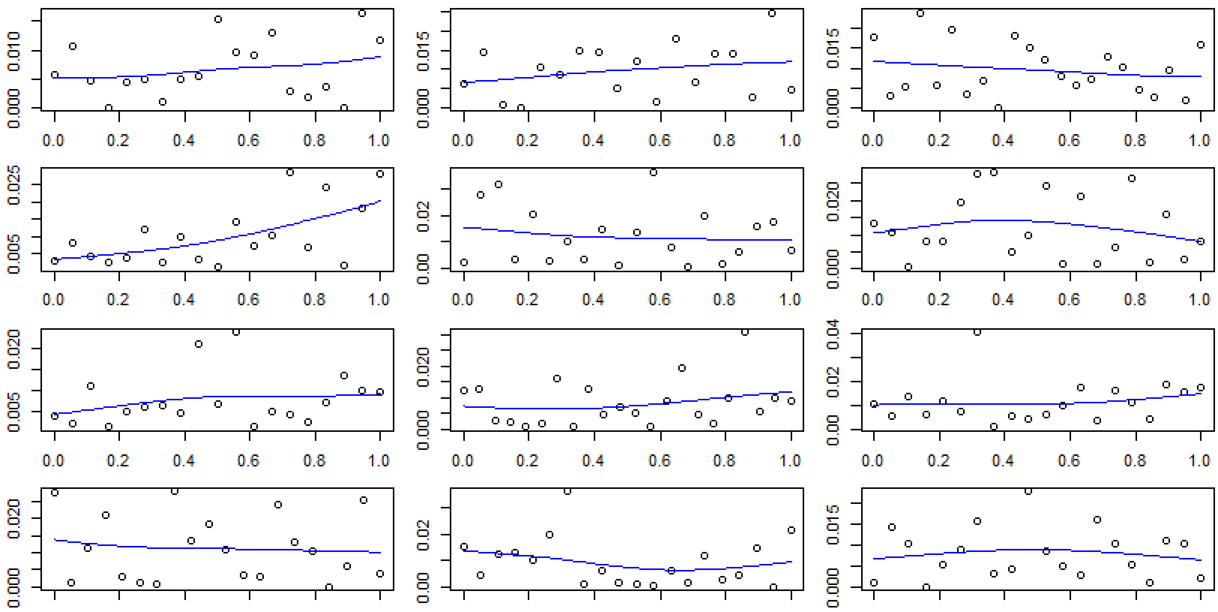

Figure 1: 2022 (12 months) volatility data points (in black circles) superposed by corresponding smoothed intra-month volatility curves (in blue lines). The vertical axes represent volatility values, and the horizontal axes represent the standardised intra-month time.

3.3 Functional Concurrent AR(1) Model with $S_{training}^{smoothed}$ as training set

The 48 smoothed intra-month volatility curves from the training set $S_{training}^{smoothed}$ are visualised in Figure 2. We observe that the green dashed curve in Figure 2 may be a potential outlier, which should be carefully dealt with. We find that this potential outlier is actually related to the stock market crash in 2020 March. A stock market crash is actually informative in a business cycle, so we decide to keep this meaningful piece of data in the training set.

Then, we use the training set $S_{training}^{smoothed}$ to train a functional concurrent AR(1) model following the method discussed in section 2.3. We obtain the estimation for the (functional) slope coefficient $\hat{\beta}_1(t)$ which is visualised in Figure 3. We observe that there is a positive correlation between this-month volatility trajectory $V_i(t)$ and last-month volatility trajectory $V_{i-1}(t)$. Further, the correlation grows stronger as t gets larger.

3.4 Functional Permutation F-Test with DJIA Volatility Data

We employ a functional permutation F-test as introduced in section 2.4 to test for the significance of the predictive relationship between last-month volatility $V_{i-1}(t)$ and this-month volatility $V_i(t)$. The test result is visualised in Figure 4. We observe that when $t > 0.2$, the predictive relationship is statistically significant at the 5% level.

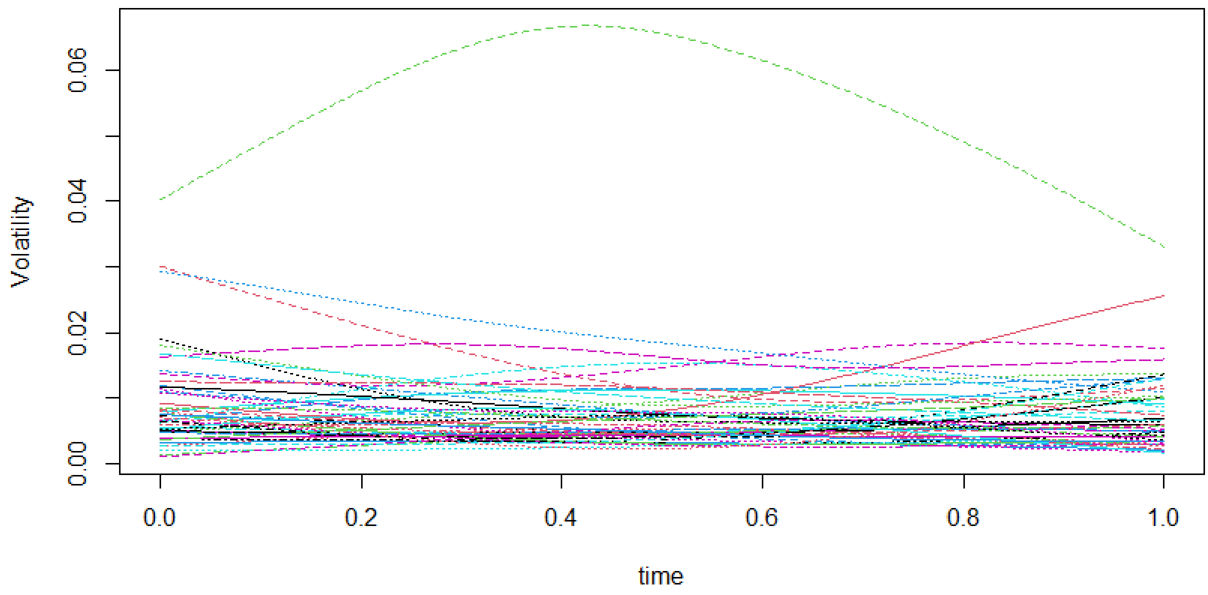

Figure 2: 48 smoothed intra-month volatility curves (2018-2021, 48 months) from the training set. The horizontal axis represents the standardised intra-month time.

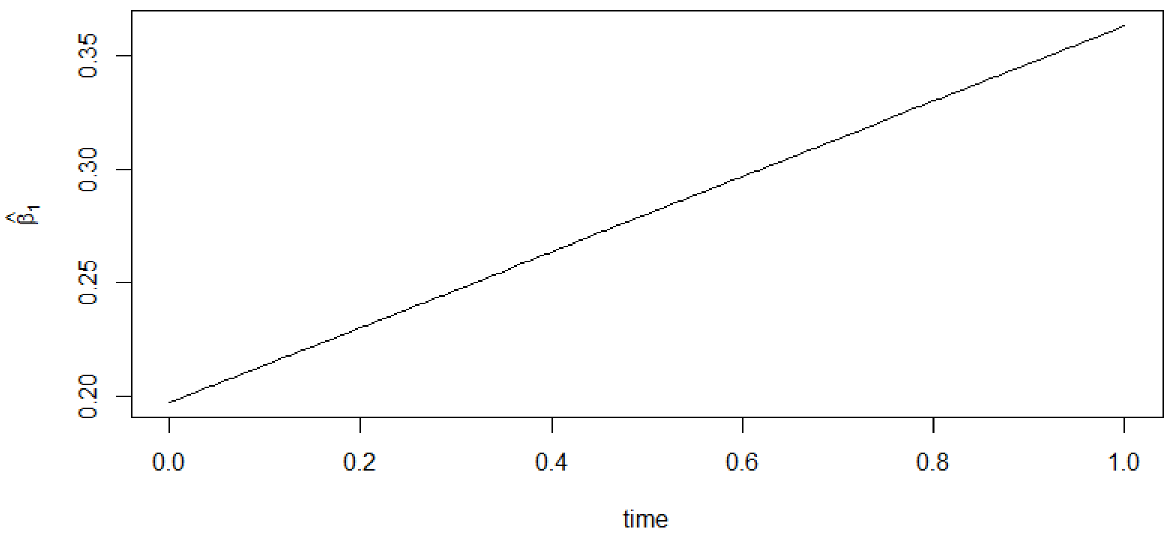

Figure 3: Estimation for the (functional) slope coefficient $\hat{\beta}_1(t)$. The horizontal axis represents the standardised intra-month time t.

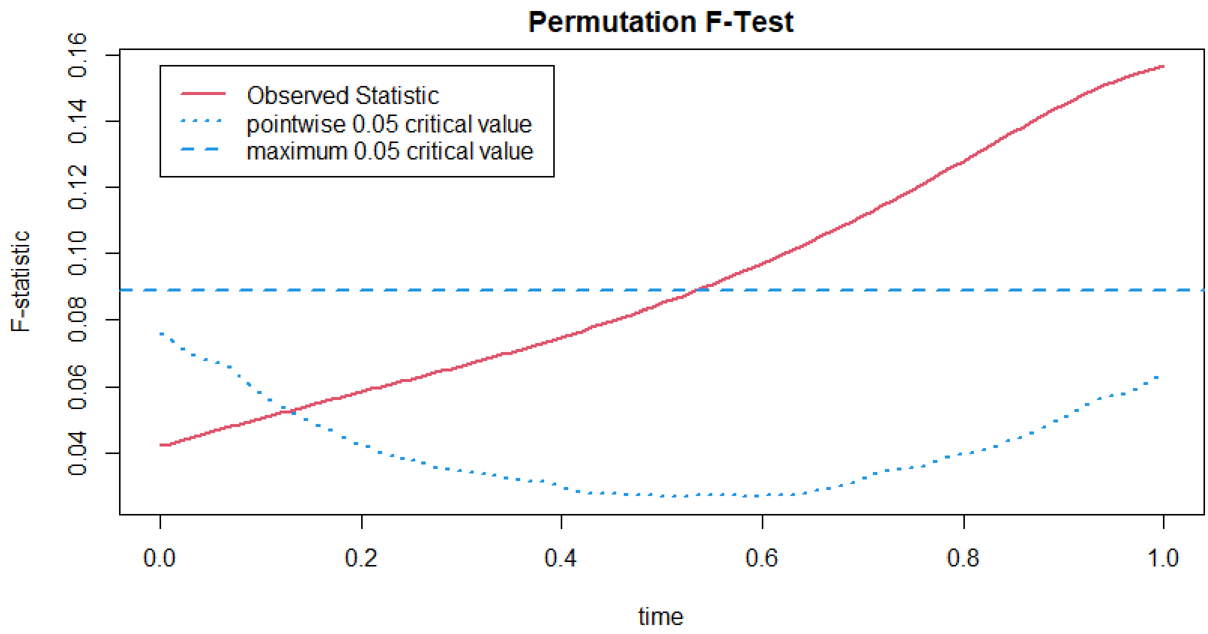

Figure 4: Functional permutation F-test for the significance of the predictive relationship between last-month volatility $V_{i-1}(t)$ and this-month volatility $V_i(t)$. The horizontal axis represents the standardised intra-month time t.

3.5 DJIA Intra-month Volatility Forecasting and Validation

We do a 12-step forecasting for 12 intra-month volatility trajectories corresponding to 12 months in year 2022. We use the model trained in section 3.3 to do forecasting. The forecasting results are visualised in Figure 5. The 12 forecasted intra-month volatility curves are shown in red lines, in contrast to the blue lines obtained by direct curve smoothing with the daily volatility data points. We can visually see that the forecasted curves and the directly smoothed curves inherit similar patterns and both give reasonable fits to the observed daily volatility data points (shown in black circles).

We would also like to validate the forecasting performance quantitatively based on the weighted mean square error (WMSE) as the goodness of fit measure discussed in section 2.5. As in Figure 6, the 12 WMSEs corresponding to the 12 forecasted intra-month volatility curves are plotted in the red line, in contrast to the 12 WMSEs of the directly smoothed curves plotted in the blue line. We observe that for short-term forecasting (e.g., the first 6 months in 2022), the forecasting performs reasonably, similar to direct curve smoothing. But for long-term forecasting (e.g., the last 6 months in 2022), the forecasting performs relatively unstable and sometimes not ideal. This may be explained by the increasing uncertainty and accumulated forecasting error in the long term. Another reason may be that our functional concurrent AR(1) model is relatively simple and not sophisticated enough, which we intend to extend and improve in future research.

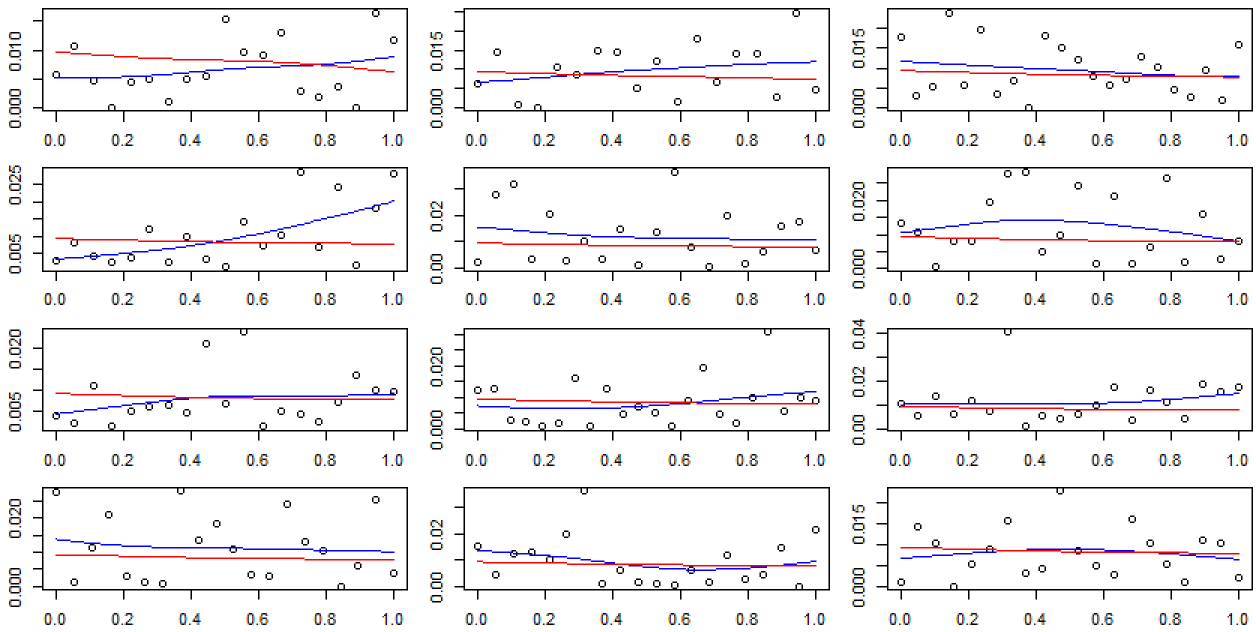

Figure 5: 2022 (12 months) forecasted intra-month volatility curves (in red lines). For contrasting purposes, daily volatility data points are shown in black circles superposed by directly smoothed intra-month volatility curves shown in blue lines. The vertical axes represent volatility values, and the horizontal axes represent the standardised intra-month time.

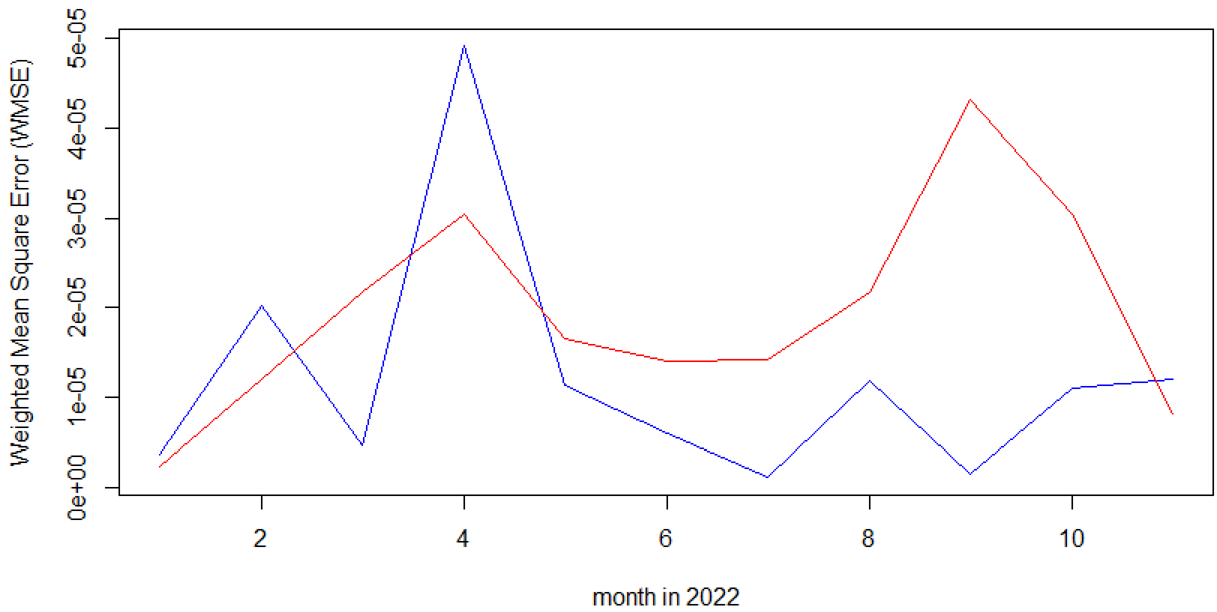

Figure 6: WMSEs for the 12 forecasted intra-month volatility curves (in red line) vs. WMSEs for the 12 directly smoothed intra-month volatility curves (in blue line). The horizontal axis represents 12 months in 2022.

4 Discussion and Conclusion

In this project, we studied intra-month stock price volatility with the help of Functional Data Analysis (FDA). Specifically, we used a functional concurrent AR(1) model to do intra-month volatility trajectory forecasting. We first introduced some essential FDA techniques and then we empirically applied them to a Dow Jones Industrial Average (DJIA) dataset. Typically, we trained a functional concurrent AR(1) model with the DJIA data and did some intra-month volatility forecasting and validation. We found that for the stock price index DJIA, there is a statistically significant predictive relationship between last-month volatility and this-month volatility. Furthermore, our functional model, or in a broader sense, FDA has provided an insightful view in stock market research and forecasting. The flexibility and information richness inherent in functional regression and forecasting are highly valuable in today's stock market, in support of rapid and continuous decision makings and relevant big data analysis.

Considering the great potential of FDA in stock market research and forecasting, we would like to point out the following possible future directions. One direction is that we would like to extend our functional model to a more sophisticated form, intending to improve forecasting accuracy and stability. For the extension, we may consider including more covariates in the model such as seasonality effects or some possible correlations between non-concurrent time points (e.g., see section 2.3). Another direction is investigating FDA further in other fields of stock market research, such as stock price prediction or high-frequency trading strategies.

Acknowledgements

I would like to thank AMSI for offering me this meaningful chance to conduct research in Mathematics and Statistics. I have learned a lot from the experience. I would also like to thank the School of Mathematics and Statistics, the University of Melbourne for the general support and guidance. Especially, I would to thank my supervisor, Professor Guoqi Qian for his insightful views, valuable suggestions and strong support along the way.

References

Alva, K, Romo, J, & Ruiz Ortega, E 2009, *Modelling intra-daily volatility by functional data analysis: an empirical application to the Spanish stock market*, Universidad Carlos III de Madrid - Working Papers.

Binder, JJ & Merges, MJ 2001, 'Stock Market Volatility and Economic Factors', *Review of Quantitative Finance and Accounting*, 17:5–26.

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

Conerly, B 2014, High Frequency Trading Explained Simply, viewed 6 January 2023. https://www.forbes.com/sites/billconerly/2014/04/14/high-frequency-trading-explained-simply/?sh=35ba50ff3da8

Craven, P & Wahba, G 1979, 'Smoothing noisy data with spline functions: Estimating the correct degree of smoothing by the method of generalized cross-validation', *Numerische Mathematik*, 31:377-403.

Huber, C, Huber, J, & Kirchler, M 2022, 'Volatility shocks and investment behavior', *Journal of Economic Behavior & Organization*, 194:56-70.

Ivrendi, M, & Guloglu, B 2012, 'Changes in Stock Price Volatility and Monetary Policy Regimes: Evidence from Asian Countries', *Emerging Markets Finance & Trade*, 48:54–70.

Macrotrends 2023, Dow Jones - 10 Year Daily Chart, viewed 19th January 2023. https://www.macrotrends.net/1358/dow-jones-industrial-average-last-10-years

Ramsay, JO, Hooker, G & Graves, S 2009, *Functional data analysis with R and MATLAB*, Springer, New York.

Ramsay, JO & Silverman, BW 2005, *Functional Data Analysis*, Springer, New York.