Forecasting high-frequency financial time series: an adaptive learning approach with the order book data

Parley Ruogu Yang*

11 Sep 2020

ABSTRACT

This paper proposes a forecast-centric adaptive learning model that engages with the past studies on the order book and high-frequency data, with applications to hypothesis testing. In line with the past literature, we produce brackets of summaries of statistics from the high-frequency bid and ask data in the CSI 300 Index Futures market and aim to forecast the one-step-ahead prices. Traditional time series issues, e.g. ARIMA order selection, stationarity, together with potential financial applications are covered in the exploratory data analysis, which pave paths to the adaptive learning model. By designing and running the learning model, we found it to perform well compared to the top fixed models, and some could improve the forecasting accuracy by being more stable and resilient to non-stationarity. Applications to hypothesis testing are shown with a rolling window, and further potential applications to finance and statistics are outlined.

JEL classification: C40, C52, C53, C58.

MSC2020 classification: 62M10, 68T05, 91B84. 91G15.

Key words: forecasting methods, statistical learning, high-frequency order book.

ACKNOWLEDGEMENT

I thank Dr Mihai Cucuringu (Department of Statistics, University of Oxford and The Alan Turing Institute) and Dr Alex Shestopaloff (School of Mathematical Sciences, Queen Mary University of London and The Alan Turing Institute) for their advice and support during the research. I thank CIFCO Guangzhou for providing the high-frequency dataset for which the empirical study can be based upon, and I am also grateful to St Anne's College, University of Oxford for its Graduate Student Research Grant, for which the cost of running computing machine is partially funded.

*Faculty of Mathematics, University of Cambridge, and Department of Statistics, University of Oxford.
Contact: ry266@cam.ac.uk

Contents

1	Introduction and Literature Review	1
2	Data	2
2.1	Data description, data cleaning and computing deployment	2
2.2	Feature generation from the order book	4
2.3	Exploratory data analysis	5
3	Methodology and Results	10
3.1	General setting	10
3.2	The fixed models: the univariate and the multivariate	11
3.3	Results from the fixed models	14
3.4	Adaptive learning model groups	17
3.5	Results from adaptive learning model groups	20
3.6	A closer look at adaptive learning results	24
4	Application to Hypothesis Testing on models	27
4.1	Method	27
4.2	Result	28
5	Discussion, Extension, and Conclusion	33
A	Additional statistical notes	35
A.1	On the ADF test	35
A.2	On MAE and MSE	35
B	Additional figures and tables	36

1 Introduction and Literature Review

Time series can be described as a sequence of observations indexed by the time, which, by the nature of it, can be separated into the past and the future. The study of predicting the future based on the past information is defined as forecasting, which is of great importance to the society — forecasting financial time series, e.g. the price of an asset, can be influential to the decisions of both the private and the public sectors. As the computerisation of financial markets develops, higher frequency of the observation on the variables are taken and can be analysed. Consequently, forecasting such a high-frequency object becomes increasingly important.

At a higher level, the statistical approach undertaken for learning the big data and obtaining better prediction has evolved in the recent decades, in both the theory (Shalev-Shwartz and Ben-David 2014) and the applications, e.g. LSTM and deep learning (Hochreiter and Schmidhuber 1997; Goodfellow, Bengio, and Courville 2016). Various methods of learning, e.g. clustering, neural networks and other synthetic models have been developed and many of which have helped to solve socio-economic problems (Aghabozorgi, Shirkhorshidi, and Wah 2015; Chakraborty and Joseph 2017). The application of statistical learning algorithms to dynamically assess forecasts has also shown a contribution to the empirical time series econometrics literature (Yang 2020).

However, due to the nature that time series dataset is indexed by time, and the fact that many structures (mathematically, such a concept is quantified by functional forms and parameters) change over time, one needs to pay particular attention while applying generically-developed learning methods to a financial time series environment (Sirignano and Cont 2018). Econometricians refer such a unique time series issue as "time-varying parameters", which could also relate to the stationarity of a model — essentially questioning the validity of the boundedness of the variables over time.¹ Recent proposals on dealing with these have been suggested by Andres and Harvey (2012) and Harvey (2013), with some empirical studies being done (Harvey and Sucarrat 2014). Additionally, adapting learning methods to improve traditional ARIMA models' forecast has also been studied empirically (Z. Li, Han, and Song 2020). In this paper, we use the classical approach of window-estimation, thereby focusing on the contemporary relationship between variables to ensure time-variability. The choice of window sizes vary, as shown later in the adaptive learning, different sizes could be preferred from time to time.

In terms of generating features (explanatory variables) to help to forecast the price, the order book data becomes particularly helpful. In this paper, the order book data consists of the quantities and prices for the best bid and ask — meaning the ones at which the asset can be traded immediately sold and bought, respectively. Statistics of these can be summarised into order flow imbalance, which synthetically involves the prices and quantities on both sides, or order imbalance, which deals with solely the quantities on both sides. The earlier has been studied by Cont, Stoikov, and Talreja (2010) and Cont, Kukanov, and Stoikov (2014) and the latter by Avellaneda, Reed, and Stoikov

¹ One may refer to Klenke (2013) for a more rigorous definition on stationarity.

(2011) and Stoikov (2014).

While more sophisticated learning models could be used, e.g. deep learning models (Sirignano and Cont 2018), we start from a traditional time series modelling and statistical learning perspective and subsequently propose learning models that can adapt to the past (thus called adaptive learning). Such a model has a better interpretability and numerous potential applications, e.g. hypothesis testing. A general reference on the foundation of model selection is from Akaike (1974) and Hastie, Tibshirani, and Friedman (2001). Additionally, penalisation plays a key role in the study of time series model selection (Cai and Wang 2014; Zbonakova, X. Li, and Härdle 2018), and functional penalisations, for instance, the MDL criterion has also shown its empirical usefulness (Rubio et al. 2007). These motivate the formation of adaptive learning proposed here.

Standard mathematical concepts and statistical notations are used in this paper. The concept of functional sets, as used in Vapnik (2000), is highlighted later in section 3 as will be used frequently. Standard time series notations are used throughout, with the main reference on stationarity issues being Banerjee et al. (1993) and Harris and Sollis (2003) and other ARIMA modellings being Harvey (1993) and Fuller (1996). Bayesian hypothesis test is adapted at the level of Koop, Poirier, and Tobias (2007), with standard frequentist test being assumed at the level of Casella and Berger (2008).

Details of the data cleaning, feature generation, and their exploratory analysis are written in section 2. Modelling and learning proposals, together with the results are presented in section 3, followed by an application to hypothesis testing in section 4 and the discussions in section 5.

In terms of the key contribution from this paper, we engage with the existing methods on feature generation and traditional ARIMA formation, then propose a forecast-centric learning model, the adaptive learning, to approach model selections and post-estimation penalisation. Such a learning model helps to confirm the reachability of certain level of accuracy of the forecasting, improves the forecasting in volatile and non-stationary markets, and can also be applied into further analogies such as on its formation and hypothesis testing.

2 Data

2.1 Data description, data cleaning and computing deployment

2.1.1 Description, time brackets and VWM

We base the statistical modelling on the intra-day price data of the CSI 300 Index Futures (hereafter called "the asset") provided by CIFCO Guangzhou. We focus the time range from 10th November 2017 to 17th April 2018, and there are 105 trading days in the range.

In a usual trading day, there are two trading sessions: one in the morning (0930—1130) and the other in the afternoon (1300—1500).² In the dataset, we expect one or two raw entries within each

² Further details can be retrieved from the exchange website: http://www.cffex.com.cn/en_new/

second, while some omissions occur throughout. In each of the raw entry, the best bid and ask data together with the latest traded price and volumes are observed. As a decision to summarise the data, we divide each of the session into 24 brackets of 5-minute slots and index them by the end time, e.g. 0935 refers to the bracket from the first second of 0930 to the last second of 0934.³ This is supported by the fact that noisiness and emptiness of the data and lack of transactions do exist. If summaries were to be made on a minute-level basis, while any brackets larger than 5 minutes would be less regarded as a high-frequency time series, as each session only has 2 hours.

A summary of statistics and a histogram of the number of observations within each bracket are available in [Table 11](#) and [Figure 18](#) respectively in the appendix. [Table 1](#) below serves as an example of translation between the brackets, thereafter "observations", the time, and in the financial environment.

Number of Brackets	12	24	48
Time	1 hour	2 hours	4 hours
Financial Remark	half trading session	1 trading session	1 trading day

Table 1: A guide between the number of observations, time, and financial meanings.

Within each bracket, we compute the Volume-Weighted-Mean (VWM) of the asset price. This is achieved by obtaining the arithmetic sum of the product of the trading volume and price in each of the raw entry, divided by the total volume. Summary of statistics is supplied in [Table 11](#) in the appendix, and a line plot of the VWM is supplied in [Figure 1](#) below.

Figure 1: Plot of the VWM of the CSI 300 Index

[CSI300IndexFutures.html](#).

³ In HHMMSS format, that is from 093001 to 093500.

2.1.2 Computing deployment

Due to the computing complexity, cloud resources with parallel computing techniques are utilised. In particular, we deploy multi-core parallel computing tasks using the machines on Google Colab and AWS, which is achieved by centralising the function and distributing the parameters we wish to compute over different CPU cores, followed by result collections individually. Codes for execution and visualisation are written in Python 3.

2.2 Feature generation from the order book

As suggested by the literature review, we generate the Order Imbalance (OIB) and the Order Flow Imbalance (OFI) for each of the raw entries as follows:

$$e_s^{OIB} = \frac{BQ_s - AQ_s}{BQ_s + AQ_s} \quad (1a)$$

$$e_s^{OFI} = BQ_s \mathbb{1}[BP_s \geq BP_{s-1}] - BQ_{s-1} \mathbb{1}[BP_s \leq BP_{s-1}] \\ + AQ_{s-1} \mathbb{1}[AP_s \leq AP_{s-1}] - AQ_s \mathbb{1}[AP_s \geq AP_{s-1}] \quad (1b)$$

where we use s as the index label for the time of the raw entry, BQ_s, AQ_s as the best bid and ask quantities respectively, and BP_s, AP_s as the best bid and ask prices respectively.

We provide a general interpretation without going deep into the theory here. For the OIB, when the bid quantity is relatively high, the OIB is more positive and vice versa if the ask quantity is relatively high. For the OFI, it can be seen as a signed contribution of the order book events to the supply or demand of the market of the asset. Say if someone buys passively through the current bid price, then $e_s = BQ_s - BQ_{s-1}$ represents the size of that order cancellation. If the bid price were to change — depending on up or down, e_s can represent the size of a price-improving order ($e_s = BQ_s$ if $BP_s > BP_{s-1}$), serving as a quantity for a rise in the demand; or the last order in the queue that was removed ($e_s = -BQ_{s-1}$ if $BP_s < BP_{s-1}$), thus a quantity for a drop in the demand. Likewise for the ask side symmetrically, where an increase in AP , for example, signifies a decrease in supply of the asset.

Within each time bracket, we need to find representable summaries of statistics to represent the behaviour of each of the two features within, technically, $\{e_s^{OIB}, e_s^{OFI} | t-1 < t(s) \leq t\}$ where $t(s)$ indicates the time bracket that entry s belongs to. Now, the mean within each bracket are the usual choice and is consistent with the intuition. In addition, we consider a p-score defined as

$$p\text{-score}_t := \Phi \left(\frac{\text{mean}(\{e_s | t-1 < t(s) \leq t\})}{\text{sd}(\{e_s | t-1 < t(s) \leq t\})} \right) \quad (2)$$

where $\Phi(\cdot)$ is the normal CDF and $\text{mean}(\cdot)$ and $\text{sd}(\cdot)$ are the mean and standard deviation of the sequence. Summary of statistics of all of these feature generated are presented in Table 11, with a line plot below in Figure 2. The benefit of having a normal transformation, as seen from the plot or summary, is that the value can be restricted into a small range (theoretically $[0, 1]$), which deals with

any potential spiky moves of the fraction, while the cost is the decrease of variance associated with the increase of stability, which is not a huge trouble as it also brings time series models a benefit of stationarity.

Figure 2: Plot of the four features generated

2.3 Exploratory data analysis

2.3.1 Initial prices of the session: exemption and dummies

Here we draw particular attention for the first and last observation of each session. Economically speaking, between the two trading sessions there could be large underlying events causing potential price movements, while the market is not open. This creates a high potential for the difference between the closing price of the previous session and the opening price of the current session to be large. We investigate these differences below.

	Day Gap	Lunch Gap	The rest
count	104	105	4830
mean	1.92	-0.44	-0.10
std	25.48	3.76	4.73
min	-134.03	-16.01	-52.51
max	79.36	6.75	32.49

Table 2: Summary of statistics for the gaps.

As shown in Table 2, the day gap, i.e. the difference between the first VWM observation in a morning's session and the last in the previous afternoon's session, is distributed much wider and have extreme values compared to the rest. This can be additionally supported by the histogram in Figure 18 in the appendix. While lunch gap, i.e. the gap between the start of the afternoon's session and the end of the morning's session, is small, to ensure consistency we exclude both of these gaps from estimation. This is done by adding dummies when the time lands at these points.

To additionally ensure the stableness of forecasting models and that it has the ability to learn the new environment within each session before making forecasts, we exclude the first 6 observations, i.e. 30 minutes, of the session from forecasting.

2.3.2 Rolling ADF tests

One crucial concern of time series is its stationarity. The approach undertaken to test the null hypothesis of unit root against the alternative hypothesis of stationarity, is via an Augmented Dickey-Fuller (ADF) test. Though, here we are interested in modelling the temporal relationships between variables, for which we are more concerned with the stationarity of the dependent variable, VWM of the price, in a short window. Hence we introduce the rolling ADF test, for which we collect the p-value, i.e. the probability of rejecting the null conditional on the null being true⁴, over time.

Let y_t be the VWM of the price at bracket t . Then, given a choice of window size w , we run an ADF test on the set $\{y_\tau\}_{\tau=t-w}^{t-1}$ for every t , and collect the result as $p_t(d; w)$ where d is the level of difference.⁵ Interpretation of the result can be made by observing $p_t(d; w)$ against a critical value, which we take as 0.05 as usual.

Starting from $d = 0$, if $p_t(0; w) < 0.05$, we conclude $\{y_\tau\}_{\tau=t-w}^{t-1}$ is stationary, else we seek for a higher order iteratively: until $p_t(d; w) < 0.05$ where we conclude $\{\Delta^d y_\tau\}_{\tau=t-w+d}^{t-1}$ is stationary while $\{\Delta^{d-1} y_\tau\}_{\tau=t-w+d-1}^{t-1}$ is not.⁶

We run this for three window sizes (12, 48, 96) and levels of difference (0,1,2) and draw, in [Figure 3](#), the line plot and histograms of the p-values for each of the combination.

The choppiness of the p-values for small-window sized data (when $w = 12$) are significantly shown, while the larger ones seem stable with occasions where $p_t(1; w) > 0.05$, meaning occasionally the 2nd level difference would be required for there to be a stationary model.

These exploratory results help to decide how the time series model should be formed, as detailed in [section 3](#).

⁴ Which is also the type I error.

⁵ Further details on the deployment of the test can be seen in [subsection A.1](#).

⁶ Δ here is the difference operator, e.g. $\Delta y_t = y_t - y_{t-1}$ and $\Delta^k y_t = \Delta^{k-1}(y_t - y_{t-1}) \forall k \geq 2$.

Figure 3: Results from rolling ADF tests for different window sizes: 12 (top two), 48 (centre two), and 96 (bottom two).

2.3.3 The SR statistics: a trading perspective

A natural extension from a financial time series model is its profitability from trading. One good model should produce a reasonable return while maintaining suitable risks. This performance can be evaluated by the Sharpe Ratio (SR). Here we explain the construction towards such a statistics and provide baseline and feature-based results.

Let P_t be the price of the asset at time t . Then the return for buying it at time t and selling it at time $t + 1$ is $\frac{P_{t+1} - P_t}{P_t}$. For each trading session after the forecasting exemption, we have 17 such opportunities, hence, given a theoretically zero-mean time series feature α_t for which the sign indicates the forecasted direction, we set the profit or loss in the trading session ($PL_s^{session}$) as

$$PL_s^{session} := \sum_{t=t_s+1}^{t_s+17} \text{sign}(\alpha_t) \frac{P_{t+1} - P_t}{P_t} \quad (3)$$

where $t_s + 1$ locates the time index to the start of the session, and accordingly $t_s + 18$ is the last observation of the session. For standard reporting on day profits or loss and further SR computation, we also produce the profit or loss in the trading day (PL_d^{day}) as

$$PL_d^{day} := PL_{s_d}^{session} + PL_{s_d+1}^{session} \quad (4)$$

where s_d locates the morning session of a trading day d . Now, the annualised SR is defined as

$$SR = \sqrt{252} \frac{\text{mean}(PL_d^{day})}{\text{sd}(PL_d^{day})} \quad (5)$$

where $\text{mean}(PL_d^{day})$ and $\text{sd}(PL_d^{day})$ stand for the mean and standard deviation of PL_d^{day} respectively.

In the baseline situation, we consider a buy-and-hold treatment, hence $PL_s^{session}$ is defined by simply buying from the start and selling at the last, thus $\frac{P_{t_s+18} - P_{t_s+1}}{P_{t_s+1}}$. Other statistics follows.

The results of these are plotted in [Figure 4](#), which clearly shows the inability for the features themselves to achieve positive returns, while the baseline also performs badly.

Figure 4: Plot of the cumulative PL_d^{day} over time

	Baseline	OFI mean	OIBmean
mean (PL_d^{day})	-0.0016	-0.0010	-0.0017
std (PL_d^{day})	0.0093	0.0092	0.0063
min (PL_d^{day})	-0.0394	-0.0397	-0.0184
max (PL_d^{day})	0.0210	0.0206	0.0197
SR	-2.71	-1.79	-4.20

Table 3: Summary of statistics of PL_d^{day} and SR

As a remark when SR is served as a performance metric later, the α_t is naturally set as $P_{t+1|t} - P_t$ where $P_{t+1|t}$ is the forecast of P_{t+1} at time t , and Equation 3 can be interpreted as the trading profit or loss if one buys whenever the next price is forecasted to raise or sell otherwise.

3 Methodology and Results

3.1 General setting

Let y_t be the one-dimensional discrete time series of interest (the dependent variable), and let $\mathbf{x}_t$ be the multi-dimensional discrete time series of features (explanatory variables). We are interested in forecasting the one-step-ahead future of the dependent variable conditional on the information up to time t , namely $y_{t+1|t}$.

In the common practice of time series, one studies the model of the underlying process and then use the model to conduct forecasting (e.g. Prado and West (2010)). Here we take a different approach: we first appreciate the conditional forecast as a value from a map that takes the information set (Φ_t) and functional parameters (θ_t, h_t) , then build models to learn the appropriate parameters based on the previous observations. Mathematically,

$$y_{t+1|t} = f(\Phi_t; \theta_t; h_t)$$

where $\Phi_t := \{(y_t, \mathbf{x}_t), (y_{t-1}, \mathbf{x}_{t-1}), \dots, (y_1, \mathbf{x}_1)\} = \{y_\tau, \mathbf{x}_\tau | \tau \in [t]\}$, $\theta_t \in \Theta(h_t)$ is the parameter to be specified in the function, and $h_t \in H$ specifies the functional form, thus determines the parameter space $\Theta(h_t)$.⁷ We manually design sensible models to construct H and sensible learning methods on θ_t and h_t to do good on reducing forecasting error — we consider the usual MSE and MAE as the performance indicators.⁸ MSE has a better theoretical foundation while MAE is more interpretable.⁹ We also consider the SR statistics which serve as an interpretable performance indicator in the context of financial time series.

At each of time t , $\mathbf{x}_t$ is four dimensional: the first and second entries are, respectively, the mean of OIB and OFI; the third and fourth entries are, respectively, the p-score of OIB and OFI.

We index each $h \in H$ by $h(\iota, w, p, d, q)$, which controls an ARIMAX(p,d,q)-type of forecasting model with w for window size and ι for the explanatory variables. For a given p, d, q , we consider a forecasting formula

$$y_{t+1|t}(h) = c_t + g_{1,t}(p, d)y_t + g_{2,t}(q)\hat{\varepsilon}_t(h) + g_{3,t}(\mathbf{x}_t) \quad (6)$$

where c_t is a constant, g_1 and g_2 are the appropriate ARIMAX operator functions: g_1 is specified by the autoregressive lag p and difference parameter d , and g_2 is specified by the moving average lag q . $\hat{\varepsilon}_t(h)$ are the residuals from the model h . $g_{3,t}(\mathbf{x}_t)$ summarises the explanatory variables' contribution to forecasting.

In the followings, we first elaborate each of the specifications of h with the associated method to pin down θ_t , thus named "fixed models", then discuss adaptive learning models where h_t can be

⁷ Here we note $\{f(\cdot; \cdot; h) | h \in H\}$ is a set of functions, this is the same notation as was used by Vapnik (2000).

⁸ See subsection A.2 for exact formulation.

⁹ See the theoretical foundation of MSE, for instance, at the Corollary 8.17 of Klenke (2013) where we view conditional expectation as projection.

time-varying by learning from the past. A general computing approach to obtain the result is shown in [algorithm 1](#).

Algorithm 1: Algorithm for obtaining the forecasts with a fixed h (fixed models)

Input: Data $\{\Phi_t\}_{t \in T}$, specification of h , desired forecasting index set T , and validation data $\{y_{t+1}\}_{t \in T}$.

Output: Forecasts $\{y_{t+1|t}(h)\}_{t \in T}$ and the performance metric.

1. For $t \in T$, repeat:
 - (a) Train parameters $\theta_t(h)$ on the windowed dataset $\Phi_t \setminus \Phi_{t-w}$, then obtain forecast $y_{t+1|t}(h) = \mathbb{E}[y_{t+1} | \theta_t(h), \Phi_t, h]$
2. Evaluate the performance metric.

3.2 The fixed models: the univariate and the multivariate

3.2.1 Univariate framework

In univariate models, the strategy to train the parameters $\theta_t^* \in \Theta(h)$ is rather classical: for a given $h(\iota, w, p, d, q)$ with $\iota \in \{0, \dots, 6\}$, we fit the following model in a w -windowed dataset:

At time $\tau \in \{t - w + 1, \dots, t\}$:

$$\Gamma_t(p)\Delta(d)y_\tau = \mu_t + \Phi_t(q)\varepsilon_\tau + \langle \beta_t(\iota), \mathbf{x}_{\tau-1} \rangle + dum_\tau; \varepsilon_\tau \sim iidN(0, \sigma_t^2) \quad (7)$$

$\Gamma_t(p), \Delta(d), \Phi_t(q)$ are the lag operator functions under an ARIMAX (p,d,q) specification with constant μ_t , e.g. $\Gamma_t(p) = (1 - \gamma_{1,t}L - \dots - \gamma_{p,t}L^p)$ where L is the lag operator, i.e. $Ly_t = y_{t-1}$. The dum_τ term dynamically adds the number of required dummies as proposed in [subsection 2.3](#).

The univariate model groups are specified by $\iota \in \{0, \dots, 6\}$. In model group 0, we set $\beta(0) := (0, 0, 0, 0)$, implying that the model is run in an ARIMA(p,d,q) fashion without explanatory variables. In model groups 1 and 2 we put $\beta(1) := (\beta_1, 0, 0, 0)$, $\beta(2) := (0, \beta_2, 0, 0)$ meaning that we solely use the mean OIB in model group 1 and the mean OFI in model group 2. In model group 3, we utilise both the mean OIB and mean OFI, thus $\beta(3) := (\beta_1, \beta_2, 0, 0)$. Likewise for model groups 4 to 6 where in 4 and 5 we consider individually each of the p-scores, then in model group 6 we combine them.

We fit the model using a Maximum Likelihood Estimation (MLE) based on the specified dataset at each time t , thus obtain the relevant parameters to implement forecasting in [Equation 6](#).

In the next paragraph we give an example to clarify the relationship between [Equation 6](#) and [Equation 7](#).

3.2.2 Univariate example

In this example we take $p = d = q = 1, \iota = 3$. Then Equation 7 becomes

$$(1 - \gamma_t L)(1 - L)y_\tau = \mu_t + (1 + \phi_t L)\varepsilon_\tau + \langle \beta_t(3), \mathbf{x}_{\tau-1} \rangle + dum_\tau; \varepsilon_\tau \sim iidN(0, \sigma_t^2) \quad (8a)$$

and we may also write $\langle \beta_t(3), \mathbf{x}_{t-1} \rangle = \beta_{t,1}\mathbf{x}_{t-1,1} + \beta_{t,2}\mathbf{x}_{t-1,2}$ in scalar form. With these specifications, we can summarise all parameters¹⁰ to estimate as $\theta_t = (\mu_t, \gamma_t, \phi_t, \beta_{t,1}, \beta_{t,2}, \sigma_t^2) \in \Theta(h) \subsetneq \mathbb{R}^6$, and in fact, by the standard time series set up we can pin down to the specified parameter region:

$$\Theta(h) = \{\theta \in \mathbb{R}^6 | \theta_2 \in (-1, 1), \theta_6 \in (0, +\infty)\}$$

Further into forecasting: once we obtained the appropriate $\theta_t^* \in \Theta(h)$, we proceed to Equation 6, which becomes

$$y_{t+1|t} = \mu_t^* + (1 + \gamma_t^*)y_t - \gamma_t^*y_{t-1} + \phi_t^*\hat{\varepsilon}_t + \langle \beta_t(3)^*, \mathbf{x}_t \rangle \quad (8b)$$

Equivalently, we can write $c_t = \mu_t^*$, $g_{1,t}(p, d) = 1 + \gamma_t^* - \gamma_t^*L$, $g_{2,t}(q) = \phi_t^*$ and $g_{3,t}(\mathbf{x}_t) = \langle \beta_t(3)^*, \mathbf{x}_t \rangle$.

3.2.3 Univariate choices of parameter

So far we explained the structure and strategy to train $\theta_t^* \in \Theta(h)$. Here we specify the choices of the model parameters. For the ARIMAX parameters, we put $p, q \in \{0, 1, 2\}$ and $d \in \{1, 2\}$, with the choice of window sizes $w \in \{12, 24, 48, 96\}$. We therefore have 72 models for each one of the seven univariate model groups, hence 504 models in total.

The reason for the window choices are from their corresponding financial meanings — as one may note from the initial data cleaning (Table 1 in particular), 12 observations refer to one trading hour while 24 refers to a session. Likewise for 48, 96 which means one and two trading days respectively. As a result, p, q may vary but rather restrictively due to the degrees of freedom, especially for smaller window sizes, hence the choice. The choice of d can be both motivated from the literature and the rolling-ADF observations done previously (Figure 3). While $d = 1$ may be sufficient, in many occasions we need $d > 1$ for stationarity purposes, hence the choices for two potential values of d .

3.2.4 Multivariate framework

In multivariate model groups, we aim for the same forecasting formula as Equation 6 but implement a vector training strategy: for a given $h(\iota, w, p, d, q)$ with $\iota \in \{7, \dots, 12\}$, we train $\theta_t^* \in \Theta(h)$ by a VARMA(p, q) on a stacked vector $S_t := (\Delta(d)y_t, \mathbf{x}_t) \in \mathbb{R}^5$. We fit the following model in a w -windowed dataset:

¹⁰ Apart from the dummies' term, which are straightforward to estimate.

At time $\tau \in \{t - w + 1, \dots, t\}$:

$$\Gamma_t(p)M(\iota)S_\tau = \boldsymbol{\mu}_t + \Phi_t(q)\boldsymbol{\varepsilon}_\tau + dum_\tau ; \boldsymbol{\varepsilon}_\tau \sim iidN(0, \Sigma_t) \quad (9)$$

Here we first note the role of $M(\iota)$: it transforms the stacked vector S_t to another which we subsequently perform VARMA on. In particular, $M(\iota) \in \{0, 1\}^{n \times 5}$ where $n \in \{2, 3\}$ is the number of parameters we plan to have. Accordingly, $\Gamma_t(p)$ and $\Phi_t(q)$ are the lag operator functions under a VARMA (p,q) specification with n dimensional variable, and the dum_τ term dynamically adds the number of required dummies as proposed in [subsection 2.3](#).

The specification on $M(\iota)$ serves in the same spirit as was the $\beta_t(\iota)$ in [Equation 7](#): it selects the relevant entries of $\mathbf{x}_t$ to interact with y_t and eventually contribute to the $g_{3,t}$ part of forecasting. For model groups 7 and 8, mean OIB and mean OFI, respectively, are the sole interaction being investigated, that is, $M(7), M(8) \in \{0, 1\}^{2 \times 5}$ and $M(7)_{1,1} = M(7)_{2,2} = 1$ with the remaining entries being zero, $M(8)_{1,1} = M(8)_{2,3} = 1$ with the rest being zero.

Similarly, training is done by MLE, and we proceed into an example.

3.2.5 Multivariate example

Consider $p = q = d = 1, \iota = 7$. Write $\widetilde{S}_\tau := M(7)S_\tau = ((1 - L)y_\tau, \mathbf{x}_{\tau,1}) \in \mathbb{R}^2$. Then [Equation 9](#) becomes¹¹

$$(I_2 - \gamma_t L)\widetilde{S}_\tau = \boldsymbol{\mu}_t + (I_2 + \phi_t L)\boldsymbol{\varepsilon}_\tau ; \boldsymbol{\varepsilon}_\tau \sim iidN(0, \Sigma_t) \quad (10a)$$

We note here $\gamma_t, \phi_t, \Sigma_t \in \mathbb{R}^{2 \times 2} \cong \mathbb{R}^4$ therefore $\theta_t = (\boldsymbol{\mu}_t, \gamma_t, \phi_t, \Sigma_t) \in \Theta(h) \subsetneq \mathbb{R}^2 \times \mathbb{R}^{2 \times 2 \times 3} \cong \mathbb{R}^{14}$

Upon obtaining the appropriate $\theta_t^* \in \Theta(h)$, forecasting proceeds:

$$y_{t+1|t} = \boldsymbol{\mu}_{t,1}^* + y_t + \gamma_{t,1,1}^*(y_t - y_{t-1}) + \gamma_{t,1,2}^*\mathbf{x}_{t,1} + \phi_{t,1,1}^*\hat{\varepsilon}_{t,1} + \phi_{t,1,2}^*\hat{\varepsilon}_{t,2} \quad (10b)$$

This is [Equation 6](#) with the specifications $c_t = \boldsymbol{\mu}_{t,1}^*$, $g_{1,t}(p, d) = 1 + \gamma_{t,1,1}^*(1 - L)$, $g_{2,t}(q) = \langle \phi_t^*, (1, 0) \rangle$ and $g_{3,t}(\mathbf{x}_t) = \gamma_{t,1,2}^*\mathbf{x}_{t,1}$.

3.2.6 Multivariate choices of parameter

For the VARMA parameters, we put $p = 1$, $q \in \{0, 1\}$, $d \in \{1, 2\}$ with $w \in \{48, 96\}$. Hence 48 models are constructed in total. One may recognise this as a more restricted choice of parameters — the choices of w are limited to the larger ones due to the degrees of freedom. Take the previous example where the parameters to estimate is equivalent to 14 dimensional, it is not realistic to be implemented when window sizes are small. For the same reason, we cap $p, q \leq 1$ while if $p = 0$ we get little meaning in the vector models, hence p is fixed at 1 and q may take one of the two values.

¹¹ Ignoring the dummy variables.

3.3 Results from the fixed models

As a summary of the results thus far, we first plot the scatter and histograms in [Figure 5](#), then, in [Table 4](#), [Table 5](#), and [Figure 6](#), we produce tables and plots for the top-performing models under the MSE ranking and the SR. [Table 12](#) in the appendix is also produced to summarise the relationship between model groups and explanatory variable(s).

For a general result, we make scatter plots and histograms for all but the outliers models — those which have an MSE greater than 100 are excluded from the plot. As can be observed from [Figure 5](#), large-window models, in general, produce lower MSE, potentially benefited from its overall stability, while outstanding small-window may also have small MSE with large SR. The linear relationship between MSE and SR is weakly negative and with many points far below or above the fitted line. This supports the discrepancy as observed later, that some models may only perform well in one of the two metrics.

Model Group	(p, d, q)	w	MSE	MAE	SR
0	(0,1,1)	96	20.97	3.34	0.13
0	(1,1,0)	96	21.03	3.35	0.69
5	(0,1,1)	96	21.22	3.36	0.99

Model Group	(p, d, q)	w	MSE	MAE	SR
5	(0,1,0)	48	22.30	3.47	5.76
2	(0,1,0)	48	22.46	3.48	4.81
5	(0,1,2)	96	27.53	3.82	4.60

Table 4: Univariate models: top 3 models ranked by MSE (upper) and SR (lower)

For top-performing models, as seen from [Table 4](#), depending on which metric we use, the "top-performing" models could vary — while models without any features (the top 2 of the upper table) perform well in MSE or MAE, their SR is rather low; with a slightly worse MSE and MAE models with features, here, in particular, the ones with either OFI mean or OFI p-score can obtain high SR, as seen from the lower table.

An interesting observation about window size may also be made — all of the models listed above are of size equal or greater than 48, similar out-performance may also be observed from the histograms of MSE in [Figure 5](#). This corroborates with the classical statistical concern on stability, as the ones with smaller window sizes may have unstable estimations which occasionally induces large errors, therefore perform badly in MSE or MAE, but not necessarily in SR.

Figure 5: Scatter plots and histograms for each of the individual models. SR is on the vertical axis and the MSE is on the horizontal axis. The red line is the linear fit between these two. The window-size 12 results are plotted in black with the window-size 24 results plotted in green on the top panel, and likewise the window-size 48 in black and 96 in green on the bottom panel. Dots represent univariate models and crosses represent multivariate models. Black and purple lines in the histogram refer to the median and mean respectively.

We take a particular notice on one 12-window-sized model with $(p, d, q) = (0, 1, 1)$ from the model group 4 — it has the fifth-highest ranking in the SR with MSE, MAE, and SR reported as 34.22, 4.22, and 4.28 respectively. This from another viewpoint shows the importance of having another performance metric — while small-window models obtain drastic forecasts from time to time, their overall ability to forecast, or at least the direction (as the SR statistic is constructed in a way that it depends on the sign of the forecast rather on the magnitude of the forecast) may still be good. In fact, when looking at vector models below, we note this phenomenon to be rather significant as shown at the top row of the lower table of Table 5. Indeed, smaller window sizes¹² cause instability, to an extent that 2 outliers of the 3600¹³ forecasts contribute largely to the bad-performing MSE and MAE.

Model Group	(p, d, q)	w	MSE	MAE	SR
10	(1,1,0)	96	21.03	3.37	1.39
11	(1,1,0)	96	21.24	3.37	1.78
8	(1,1,0)	96	21.26	3.38	1.99

Model Group	(p, d, q)	w	MSE	MAE	SR
11	(1,1,1)	48	1603.69	4.42	3.24
8	(1,1,0)	48	21.95	3.43	3.15
8	(1,1,1)	48	22.97	3.51	2.73

Table 5: Multivariate models: the top 3 models ranked by MSE (upper) and SR (lower)

In terms of the top-performing vector models ranked by MSE, it is close to the univariate results with a slightly higher SR. While the top-performing vector models ranked by SR does not outperform the ones from univariate groups. This provides evidence that vector models do not perform outstandingly well in the context of one-step-ahead forecasting and windowed estimation.

We also plot a cumulative PL of each of the top models in Figure 6 along with the aforementioned one from the 12-window-sized and the baseline. As can be seen, the univariate models can perform better than the vector models, while the large-sized univariate model, i.e. the one with the lowest MSE, has a modest level of cumulative PL throughout. This could be explained by its ability to produce low errors, benefited from its stability from larger degrees of freedom, but not able to catch the time-varying change of the underlying parameters, hence the ability to forecast the sign and make trading profits is rather low.

¹² Here we note that $w = 48$ is relatively small in the context of vector models, due to the dimension of parameters it needs to estimate. One may observe from Figure 5 that indeed the distribution of vector models for $w = 48$ is much wider compared to the same windowed univariate models, and several extreme points exist.

¹³ As checked in details of their distribution.

Figure 6: Plot of the cumulative PL_d^{day} over time (Note: the legend is coded by model group number, window size, and p, d, q values.)

3.4 Adaptive learning model groups

3.4.1 Principle of learning

In the above twelve model groups, we fix h_t constant throughout. Here, we consider an adaptive method to learn h_t . We also notice that, in the $h(11, 48, 1, 1, 1)$ for example, some drastic outliers could worsen the performance, especially if evaluated by the MSE. Hence we first shrink the functional set H by assigning certain error and outlier handling ability as below:

$$\tilde{H}_t = \left\{ h \in H \mid y_{t+1|t}(h) \text{ is positive and } \left| \frac{y_{t+1|t}(h)}{y_t} - 1 \right| < 0.05 \right\} \quad (11)$$

This makes sure the set of functions we are selecting from are not outliers. Now, let time index W be large enough so we can run all the previous model groups and obtain a suitable amount of forecasting error $\epsilon_{\tau+1}(h)$ from each of the model h , which is defined as follows. At any $t \geq W$, we have access to models in the group 0 to 12, which produce forecasts $y_{\tau+1|\tau}(h)$ for each of the $h \in H$, $t - 1 \geq \tau \geq W$. We write $\epsilon_{\tau+1}(h) := y_{\tau+1} - y_{\tau+1|\tau}(h)$ and the adaptive learning aims to learn from the errors available up to time t , together with other information available, to make a decision on the model to employ at time t .

A general strategy to train h_t at time t is to construct a loss function $\ell(\Phi_t, H) = \ell(\Phi_t, h, H \setminus \{h\})$ and solve the appropriate minimisation problem:

$$h_t^* = \arg \min_{h \in \tilde{H}_t} \ell(\Phi_t, h, H \setminus \{h\}) \quad (12)$$

The associated computing procedure is supplied in [algorithm 2](#).

Algorithm 2: Algorithm for obtaining the forecasts with a time-varying h_t

Input: Data $\{\Phi_t\}_{t \in T}$, functional sets H , specification of the loss function $\ell(\cdot, \cdot)$, desired forecasting index set T , and validation data $\{y_{t+1}\}_{t \in T}$.

Output: Forecasts $\{y_{t+1|t}(h_t^*)\}_{t \in T}$ with the associated functions $\{h_t^*\}_{t \in T}$, and the performance metric.

1. For $t \in T$, repeat:
 - (a) Produce $\tilde{H}_t$ according to Equation 11.
 - (b) Evaluate and execute the minimisation given by Equation 12. Then get h_t^* with $y_{t+1|t}(h_t^*) = \mathbb{E}[y_{t+1}|\theta_t(h_t^*), \Phi_t, h_t^*]$.
2. Evaluate the performance metric.

In the followings, we consider two groups of specifications of the loss function, with the first one motivated from Yang (2020).

3.4.2 Adaptive learning on errors

In model group 13, we construct ¹⁴

$$\ell(\Phi_t, h, H \setminus \{h\}) = L^{global}(\{\epsilon_\tau\}_{\tau=W}^t) = \sum_{t-47 \leq \tau \leq t} \lambda^{t-\tau} L^{local}(|\epsilon_\tau|) \quad (13)$$

The above line defines the loss by solely focusing on the errors in the past 48 observations, which are diminishing geometrically at a rate $\lambda \in (0, 1]$. The local loss function $L^{local}(\cdot)$ is specified as a continuously differentiable combination between zero, square loss, and absolute loss, with constants C_1, C_2 where $0 \leq C_1 \leq C_2 \leq \infty$: ¹⁵

$$L^{local}(x; C_1, C_2) = \begin{cases} (C_2 - C_1)x + \frac{C_1^2 - C_2^2}{2} & \text{if } x > C_2 \\ \frac{(x - C_1)^2}{2} & \text{if } C_2 \geq x > C_1 \\ 0 & \text{otherwise} \end{cases} \quad (14)$$

Benefited from the high-frequency dataset, differently from Yang (2020), we here place a time-varying constants at which the losses switch — we trial different quantiles of $\{|\varepsilon_t(h)| | h \in H\}$ to set $C_1(t), C_2(t)$. E.g. $C_1(t)$ can be the 25% quantile of $\{|\varepsilon_t(h)| | h \in H\}$ and $C_2(t)$ can be 50% or 75% quantile of $\{|\varepsilon_t(h)| | h \in H\}$. Note here that the quantiles have the benefit of outliers-resilient as it relates to the distribution rather than expectation of the set. The time-varying parameter here assists the local penalisation to be done in a time-varying manner and thus the minimisation process.

As to the selection of parameters, we consider $\lambda \in \{0.8, 0.85, 0.9, 0.95, 0.99, 1\}$ and 25%, 50%, 75%

¹⁴ Now and thereafter, we abbreviate $\epsilon_\tau(h)$ as ϵ_τ where h is clearly emphasised on the left hand side of the equation.

¹⁵ Note also that if $C_1 = 0$ and $C_2 = \infty$, the local loss becomes proportional to the MSE contribution, similar to the ones proposed by Yang (2020).

quantiles for the C_1, C_2 . Some interpretation on the λ can be made based on its logarithmic and exponential properties — the half period of 0.95 for example, is around 13.5 thus having $\lambda = 0.95$ essentially reviews the past 14 observations with little role being played by the further ones, while $\lambda = 0.8$ is more extreme, as its half period is only about 3.1. A general table is supplied in [Table 13](#).

3.4.3 Adaptive learning with functional awards and penalties

One concern about purely focusing on the forecasting error is the potential to misfit as the functional form plays a role in the degree of freedom and may also be of importance when selecting which model to adapt as the most appropriate one for h_t^* . The classical approach that adds penalisation on the model selection criteria is from Akaike (1974), and later the Lasso methods, in particular, the fused lasso (Tibshirani et al. 2005). Though we are different from the previous methods as we focus on the out-of-sample loss rather than the in-sample likelihood.

Here, in model group 14, we take a more time-varying approach: we consider a penalisation or reward, depending on the particular design, between the function of concern h and the previous choice h_{t-1}^* , write as $D(h, h_{t-1}^*)$.

$$\ell(\Phi_t, h, H \setminus \{h\}) = L^{global}(\{\epsilon_\tau\}_{\tau=W}^t) + D(h, h_{t-1}^*) \quad (15)$$

In the followings, we first consider, in type-1, purely penalising the difference in each of the variables that form h , i.e. the time series parameters and the window sizes.

In type-1 we design

$$D(h, h_{t-1}^*; C_3, C_4) = C_3|p + d + q - p_{t-1}^* - d_{t-1}^* - q_{t-1}^*| + C_4|w - w_{t-1}^*| \quad (16a)$$

where $C_3, C_4 > 0$.

As specified, we pool the p, d, q together and, due to the size difference, treat w separately. The fact we pool p, d, q together can be appreciated as the change in the complexity, in particular, number of lags and differences in total. Another potential way is to separate them and penalise the change one by one.¹⁶ What is also interesting to consider is the fact that we may still have a preference to large-window sized models due to their stability, hence an additional reward can be made to encourage switches into smaller window-sized models and differenced orders, hence the followings for type-2 and 3:

$$D(h, h_{t-1}^*; C_3, C_4, C_5) = C_3|p + q - p_{t-1}^* - q_{t-1}^*| + C_4(48 - w) + C_5|d - d_{t-1}^*| \quad (16b)$$

¹⁶ For instance, $C_{3,1}|p - p_{t-1}^*| + C_{3,2}|d - d_{t-1}^*| + C_{3,3}|q - q_{t-1}^*|$.

where $C_3 > 0$, $C_4, C_5 < 0$ in type 2. Now, in type-3 we put

$$D(h, h_{t-1}^*; C_3, C_4, C_5) = C_3|p + q - p_{t-1}^* - q_{t-1}^*| \mathbb{1}[d = d^*] + C_4(48 - w) + C_5|d - d_{t-1}^*| \quad (16c)$$

where $C_3 > 0$, $C_4, C_5 < 0$.

Both of the specifications have a strictly increasing reward for smaller window sizes as small-window sized ones get a more negative value on the $C_4(48 - w)$ term, which proceed into a beneficial status at the minimisation stage. The difference between the type-2 and 3 is that in type-3, penalisation on the p, q terms are only applied if d stayed the same — this paves path for potential switches in d , which may be desirable when instability breaks out and later finished.

We inherit the parametrisation in model group 13, while for the C_3, C_4, C_5 , we consider the following time-varying parametrisation. Let

$$l_t^* := \min_{h \in \tilde{H}_t} L^{global}(\{\epsilon_\tau\}_{\tau=W}^t)$$

and we set $C_3(t), C_4(t), C_5(t) \propto l_t^*$. In particular, for type-1, we put

$$C_3(t) = \frac{1}{10}l_t^*, \quad C_4 = \frac{1}{168}l_t^*$$

and for type-2 and 3,

$$C_3(t) = \frac{1}{8}l_t^*, \quad C_4(t) = -\frac{1}{72}l_t^*, \quad C_5(t) = -\frac{1}{2}l_t^*$$

The rationale behind these particular fractions is that we aim to control the maximum loss being added from each term of the $D(h, h_{t-1}^*)$ to be half of l_t^* , and likewise for the magnitude of the rewards in the type-2 and 3.

3.5 Results from adaptive learning model groups

Model Group	type	λ	(C_1, C_2)	MSE	MAE	SR
14	1	1	(50%,75%)	22.05	3.42	1.54
13		1	(50%,75%)	22.14	3.42	2.07
14	2	0.99	(50%,75%)	22.32	3.45	1.15
14	3	0.99	(50%,75%)	22.38	3.45	1.37

Model Group	type	λ	(C_1, C_2)	MSE	MAE	SR
14	1	0.85	(50%,75%)	22.52	3.47	2.90
14	3	0.8	(25%,50%)	33.14	4.15	2.63
13		0.85	(50%,75%)	22.51	3.47	2.56
14	2	0.8	(25%,50%)	28.33	3.85	2.41

Table 6: Top models from each model group-type, ranked by MSE (upper) and SR (lower)

For each of the type in model group 14 and throughout the model group 13, we select the best model ranked in either MSE or SR and list them above in Table 6. Compared against Table 4, we see the adaptively learnt models may achieve similar, though no better results, compared to the ones in the fixed models, both when measured in MSE and SR. This motivates the later subsection where we look closer into some of the periods that adaptively learnt models outperform the best one from the fixed model groups.

Figure 7: Plot of the cumulative PL_d^{day} over time

To have a more financial comparison, we observe in Figure 7, the cumulative PL plotted first within themselves and then with the top ones from the fixed model groups. Compared to the best model measured in MSE, which performs poorly in the SR, the adaptively learnt models obtain higher cumulative returns while not high enough compared to the ones which are the best individual models measured in SR.

Model code	$100 \times \text{mean}$	$100 \times \text{std}$	$100 \times \text{min}$	$100 \times \text{max}$	SR
M5_48_PDQ010	0.28	0.74	-1.02	3.72	5.96
M4_12_PDQ011	0.20	0.71	-1.55	2.34	4.38
M11_48_PDQ111	0.15	0.73	-1.17	2.63	3.29
MG14_50+75_type-1_085	0.13	0.71	-1.10	2.70	2.89
MG14_25+50_type-3_08	0.12	0.71	-1.42	2.30	2.61
MG13_50+75_1	0.09	0.73	-1.47	2.48	2.05
MG14_50+75_type-1_1	0.07	0.73	-1.16	2.48	1.53
M10_96_PDQ110	0.07	0.79	-1.26	3.42	1.46
M0_96_PDQ011	0.03	0.70	-1.74	1.95	0.58
Baseline	-0.17	0.95	-3.94	2.10	-2.76

Note about the number of days: after some data for the initialisation of models, we have 99 trading days starting from 20th November 2017, which explains the small difference between the baseline data here and the ones in [Table 3](#).

Note about the model code: for fixed models, they are coded by model group, window size, and p, d, q parameters; for the learning models, they are coded by model group, C_1, C_2 in percentage, type number (if applies), and λ .

Table 7: Summary of statistics of PL_d^{day} (columns 1 to 4, within which all entries are multiplied by 100 for cleanness) and SR (column 5)

Another financial comparison can be made from [Table 7](#), that the standard deviation, min, and max of the daily return from the adaptive learning models are all in line with the top-performing fixed models, which intuitively explains the stability of these adaptively learnt models.

The benefit of adaptive learning models is one could look into the formation of each of the models. As shown in [Figure 8](#), we may see the preferences of the model parameters through the period — univariate models encapsulate a majority, $d = 2$ is occasionally visited, and there is a good blend of usage of explanatory variables and the other time series parameters.

The choices of window sizes, depending on the design, may vary largely — in particular, M4 has a strong preference towards the smaller ones and, in terms of explanatory variables, it prefers to use none of them.

In the following subsection, we take a closer look at periods when adaptive learning results outperform the fixed models.

Model number assigned for Figure 8	Model Group	type	λ	(C_1, C_2)
M1	13		1	(50%,75%)
M2	14	1	1	(50%,75%)
M3	14	1	0.85	(50%,75%)
M4	14	3	0.8	(25%,50%)

Table 8: Model number assigned for Figure 8

Figure 8: Plots of the formation of the four model groups: variables (top left), window sizes (top right), p, d, q parameters and their training method (bottom)

3.6 A closer look at adaptive learning results

We first observe a period when prices are volatile and most of the models obtain large forecasting errors. We zoom into a five trading day period starting from 8th February 2018. As plotted in the top of [Figure 9](#), there is a large drop and subsequently big fluctuations around. We label M_0 as the best-performing individual model in terms of MSE, which is the sized 96, $(p, d, q) = (0, 1, 1)$ model with no explanatory variable, and as we see from the second plot of [Figure 9](#), the forecasting errors can be spiky and occasionally large, contributing a large MAE and MSE on average, as tabled in [Table 9](#).

We observe from two outstanding models from the model group 14 — labelled as M_1 , the type-3 with $\lambda = 1$ and $(C_1, C_2) = (50\%, 75\%)$ and labelled as M_2 , the type-2 with $\lambda = 0.9$, $(C_1, C_2) = (50\%, 75\%)$ while $C_5(t)$ adjusted to $-l_t^*$.

Model and error	$M_0 (\epsilon_t)$	$M_1 (\epsilon_t)$	$M_2 (\epsilon_t)$	$M_0 (\epsilon_t^2)$	$M_1 (\epsilon_t^2)$	$M_2 (\epsilon_t^2)$
mean	6.30	6.14	6.08	83.99	83.44	81.94
min	0.05	0.03	0.03	0.00	0.00	0.00
max	51.81	51.52	51.52	2684.70	2654.32	2654.32

Table 9: Summary of the absolute forecasting errors (columns 1 to 3) and squared forecasting errors (columns 4 to 6)

As observed from [Table 9](#), the MAE and MSE in this period are generally high, while the adaptively learnt models have reduced them to a certain extent. The second plot of [Figure 9](#) plots the absolute forecasting error of M_0 and M_2 , and the third row of plots of [Figure 9](#) show the level of improvement (if positive) or worsening (if negative) from M_0 to M_1 on the left, and from M_0 to M_2 on the right.

The selection of window sizes may also explain the source of improvement — shorter ones are selected during a few periods when the prices are relatively unstable, and the relevant p, q parameters are changed throughout. These details can be further seen from the bottom of [Figure 9](#). Among the M_1 and M_2 , some difference can also be observed in the variable selection and training method, though most of the time no explanatory variable and univariate training methods are preferred.

The trend that adaptive learning models seem to perform better in the non-stationary part of the data motivates another review on the performance of adaptive learning models. Here we consider the non-stationary data as observed in the [subsection 2.3](#) — there were time when, at the level of window-sized 12 that we reject the null hypothesis of the ADF tests at 5% significance level for both 0-diff, 1-diff, and 2-diff, meaning that y_t is not stationary even after twice the difference. There are 152 observations of this nature that intersect with the forecasting set, and below in [Table 10](#), we provide a summary of statistics for the forecasting errors.

Figure 9: The first plot: the price during the 5 trading-day period; the second plot: plot of absolute forecasting errors of M_0 and M_2 ; the bottom four plots: error comparison and model formations of M_1 (left) and M_2 (right).

Model and error	M_0 ($ \epsilon_t $)	M_1 ($ \epsilon_t $)	M_0 (ϵ_t^2)	M_1 (ϵ_t^2)
mean	3.30	3.39	26.28	24.73
min	0.01	0.05	0.00	0.00
max	24.69	22.86	609.76	522.41

Table 10: Summary of the absolute forecasting errors (columns 1 to 2) and squared forecasting errors (columns 3 to 4)

In Table 10, we see the best-fixed model (M_0) to have moderate MAE and MSE while the adaptively learnt one (M_1) has slightly smaller MSE benefited from its avoidance from the large errors. The histogram in Figure 10 also supports such evidence — errors from M_1 have less distribution on the right tail.

Figure 10: Histograms of the absolute errors (left) and the squared errors in $\log_{10}$ scale (right)

The formation of the adaptively learnt model presented here can be found below in Figure 11: it has a high proportion of model group 0, meaning there are a vast majority of models being purely ARIMA without explanatory variables, while the window sizes tend to be the smaller ones and avoids the 96. The difference order is usually taken at 1, similar to the others observed previously.

Figure 11: Histograms of the window sizes (left), difference order (centre), and model group (right)

4 Application to Hypothesis Testing on models

Having had adaptively learnt models combining different explanatory variables, training methods, ARIMA parameters, and window sizes, one may want to test statistically the significance of certain functional classes within the selection over certain periods. Here we provide two approaches: the simple Bayesian framework where we set the prior to be proportional to the inverse of the size of the hypothesis class and updates the likelihood by simple counting, and a frequentist framework whereby binomial distribution can be assumed and thus p-values can be produced.

4.1 Method

4.1.1 A simple Bayesian hypothesis testing

For a subset of functions $H^1 \subsetneq H^0 \subseteq H$, consider the following hypothesis:

H_0 : Functions in the set H^1 have the same or lower chance of being selected than the ones in $H^0 \setminus H^1$.
 H_A : Functions in the set H^1 have a higher chance of being selected than the ones in $H^0 \setminus H^1$.

Note that here, "being selected" refers to that a particular function in the functional set H^0 or H^1 being selected by the adaptive learning model.

We consider a Bayesian framework to test the above hypothesis: write $\pi(H_0)$ and $\pi(H_A)$ as the prior for H_0 and H_A , respectively, and $\pi(y_t|H_0)$, $\pi(y_t|H_A)$ as the (conditional) likelihood for a testing observation value y_t based on the hypothesis H_0 and H_A respectively.

We define the following probabilities:

$$\pi(H_0) := \frac{|H^1|}{|H^0|} \quad (17)$$

$$\pi(y_t|H_0) := \mathbb{1}[\text{Function in } H^0 \setminus H^1 \text{ is selected}] = \mathbb{1}[h_t^* \in H^0 \setminus H^1] \quad (18)$$

$$\pi(y_t|H_A) := \mathbb{1}[\text{Function in } H^1 \text{ is selected}] = \mathbb{1}[h_t^* \in H^1] \quad (19)$$

And as a result, given a period of data $\mathbf{y} := \{y_t\}_{t=T_0}^{T_1}$, the Bayes factor for H_A is computed as:

$$B_{A0}(\mathbf{y}) := \frac{\pi(H_A|\mathbf{y})}{\pi(H_0|\mathbf{y})} \quad (20a)$$

$$= \frac{\pi(H_A) \times \pi(\mathbf{y}|H_A)}{\pi(H_0) \times \pi(\mathbf{y}|H_0)} \quad (20b)$$

$$= \frac{|H^0 \setminus H^1|}{|H^1|} \times \frac{\sum_{t=T_0}^{T_1} \mathbb{1}[h_t^* \in H^1]}{\sum_{t=T_0}^{T_1} \mathbb{1}[h_t^* \in H^0 \setminus H^1]} \quad (20c)$$

We note from this setting, that if the true underlying process was H_0 , the Bayes factor tends to 1 or lower, and otherwise higher than 1. In case the fraction in Equation 20c contains 0 in

the denominator, we assign the Bayes factor with infinity and assign a high value in the plot for illustration.

4.1.2 A frequentist hypothesis testing

Another way to test the hypothesis, or a more common frequentist approach, can be done by simply assuming, in case of H_0 , a model in H^1 is chosen with probability $\frac{|H^1|}{|H^0|}$ as we assume no better performance under the null.

Suppose we have a period of test data $\mathbf{y} := \{y_t\}_{t=T_0}^{T_1}$, then write n_1 as the number of t such that $h_t^* \in H^1$.¹⁷ Under null hypothesis, we have $n_1 \sim \text{Bin}(T_1 - T_0 + 1, \frac{|H^1|}{|H^0|})$ as there are $T_1 - T_0 + 1$ number of observations each with a likelihood of at most $\frac{|H^1|}{|H^0|}$ to be chosen.¹⁸ We can therefore set the relevant critical value for the hypothesis testing, as well as the p-value. As usual, a close-to-zero (usually set as less than 0.05) p-value indicates a rejection of null in favour of the alternative H_A .

4.2 Result

We use the above approach to test seven hypotheses, each over a period of five trading days, within which there are 180 forecasting samples.¹⁹ We use three adaptive learning models to test simultaneously: we label M_1 as the one with $\lambda = 1$ from model group 13, M_2 as the one with $\lambda = 0.8$ from model group 13, and M_3 as the one with $\lambda = 1$ from model group 14 type 1. All of them are with $(C_1, C_2) = (50\%, 75\%)$. The choices of M_1 and M_3 are because of their good performance in MSE (see Table 6) while M_2 is picked as a representative of low λ — which corresponds to a shorter-memory selection, aiding statistical conclusion here, though perform relatively badly in MSE ranking. Also, it is worth taking note that M_3 may be slightly questionable while performing the frequentist test, as the set-ups from the model group 14 could lead to highly-correlated model choices due to the term $D(h, h_{t-1}^*)$, which contradicts with the underlying assumption of binomial distribution in the null hypothesis.

The first hypothesis testing is to test the window size: whether it is 96 or not. Hence the functional set is $H^1 = \{h(\iota, w, p, d, q) | w = 96\}$ and $H^0 = H$. As shown in Figure 12, we see this to be significant on many days for both tests — a high Bayes factor can be observed in all models, for all periods but one. Likewise, a close to, if not 0, p-value can also be observed for almost all periods. Such a frequent rejection of the null hypothesis means that the adaptive learning models still have a high reliance on the large-windowed models.

¹⁷ That is, $n_1 = \sum_{t=T_0}^{T_1} \mathbb{1}[h_t^* \in H^1]$.

¹⁸ More precisely, here we follow the traditional step to obtain a Binomial distribution from $T_1 - T_0 + 1$ number of iid Bernoulli with probability being the set as the proportion.

¹⁹ When required later, a higher-frequency can also be made, for which the results are more spiky due to the small sample.

Figure 12: Bayesian testing result (left) and frequentist testing result (right)

We now consider a hypothesis about whether the model group is zero. Hence the functional set is $H^1 = \{h(\iota, w, p, d, q) | \iota = 0\}$ and $H^0 = H$. We see from Figure 13 that the significance is high in most periods. Hence we could conclude these adaptive learning models to have a high occupation of model group zero, therefore the use of explanatory variable could be low, in many periods.

Figure 13: Bayesian testing result (left) and frequentist testing result (right)

We question about the training method — whether the parameters are trained in univariate or vector models. Here we have $H^1 = \{h(\iota, w, p, d, q) | \iota \leq 6\}$ and $H^0 = H$. From the previous examples, the univariate choice may be intuitively true, as multivariate models are rarely used, and from the testing results, such an intuition is confirmed. Indeed, as shown in Figure 14, in both Bayesian and the frequentist, significance can be shown in all periods, and the magnitude of the Bayes factor is also

gigantic.

Figure 14: Bayesian testing result (left) and frequentist testing result (right)

Now, we wonder if the significance of small-window model exists, compared to the already-tested significant large-window models (from Figure 12). In this case, we have $H^1 = \{h(\iota, w, p, d, q) | w = 12\}$ and $H^0 = \{h(\iota, w, p, d, q) | w \in \{12, 96\}\}$. The p-values are almost all 1, and the Bayes factors, as shown in the left panel of Figure 15, is rather low, meaning the small-windows model are not picked significantly compared to the large one. We also attempt to "zoom-in" by using a 1-day period, for which the Bayes factors are plotted in the right panel. The p-values are mostly 1 or close-to 1, while the Bayes factors, as can be seen, are somewhat spiky — this can be due to the way the Bayes factor is constructed, but it suggests some instability as the value is spiky, for most models. This means, despite the small-window models are non-significant at a larger scale, they may occasionally play a part, as suggested by the Bayes factors on a 1-day period.

Likewise, in the followings we test the significance of certain choices of variables compared to a wider set, in particular, we first consider the sole choice of OFI mean compared against its combination with OIB mean or solely the OIB mean, which leads to $H^1 = \{h(\iota, w, p, d, q) | \iota \in \{2, 8\}\}$ and $H^0 = \{h(\iota, w, p, d, q) | \iota \in \{1, 2, 3, 7, 8, 9\}\}$. Test results, as shown in Figure 16, suggest no dominance by the OFI mean for all but one period from a frequentist viewpoint, while the Bayes factor is occasionally large, meaning significance may exist for some periods. This is due to the fact that the set by null hypothesis (H^0) may not occupy a large amount of functions being chosen,²⁰ contributing to a spiky and potentially large Bayes factor by occasion.

²⁰ As shown in Figure 13, the model without explanatory variable is of high significance.

Figure 15: Bayesian testing result in a 5-day period (left) and in an 1-day period (right)
Note on the right panel: infinity is plotted as 3 here.

Figure 16: Bayesian testing result (left) and frequentist testing result (right)
Note on the left panel: infinity is plotted as 220 here, with a logarithm scale ($\log_{10}$) being used.

We now run the same test for OFI p-score compared against its combination with OFI p-score or solely the OIB p-score. Hence $H^1 = \{h(\iota, w, p, d, q) | \iota \in \{5, 11\}\}$ and $H^0 = \{h(\iota, w, p, d, q) | \iota \in \{4, 5, 6, 10, 11, 12\}\}$. As shown in Figure 19 in the appendix, the significance is not huge although the Bayes factor may be occasionally high — this, as explained in the previous part, can be purely due to the lack of samples occupied by H^0 .

Finally, we concern about the choice of the difference order, thus $H^1 = \{h(\iota, w, p, d, q) | d = 2\}$ and $H^0 = H$. As shown in Figure 17, while testing this in a 5-day period, one may easily conclude the

insignificance of H^1 due to its consistently low Bayes factor and likewise the constantly almost-unity in the p-value; once zoomed into the 1-day period, the result becomes spiky in both the Bayesian and the frequentist tests, implying that the significance of 2nd difference being chosen is occasionally high.

Figure 17: Bayesian testing result (left) and frequentist testing result (right) with 5-day period (up) and 1-day period (down)

Note on the lower left panel: infinity is plotted as 18 here.

As a conclusion from these testing, we see a significant component of large-window sized models and models without explanatory variables being selected by the top learning models, while instability and spikiness occur, in every aspect of the model choices, especially as we zoom into a 1-day period. This is because occasionally, statistical significance can be found for 2nd difference, small-window models, as well as groups of some explanatory variables.

5 Discussion, Extension, and Conclusion

As a general discussion about the adaptive learning, we reflect first from the statistical intuition: what is h_t and why do we care to learn h_t ? The ever-changing structure, as being frequently studied, requires certain awareness by the model on the time-variability, not only on the parameters, but also the functional forms. Classical treatments on time series, e.g. Harvey (1989) and Harvey (1993) offer the chance of parameter variation, while modern deep learning models, e.g. Sirignano and Cont (2018) can help on learning the h_t .

However, some issue may occur in the learning on h_t — interpretability is one, and ultimately the design could be questionable. Here we design a model-adapting and time-adaptive learning regime, to offer higher interpretability and allow testing to be undertaken. Statistical conclusions may also be drawn from the adaptive learning models, for example, the explanatory variables may not be of good use for many occasions, as concluded in the testing.

The design, particularly Equation 12 and Equation 15, allows H to be potentially infinite. For instance, one could set H to contain infinitely many p but set the move to be at most 1 from each time, i.e. $D(h, h_{t-1}^*)$ takes infinity if $|p - p_{t-1}^*| > 1$. This allows the model to be theoretically more variable and contributes to the ultimate learning on h_t . Another more computationally expensive, but also important extension is to dynamically learn the "hyper-parameter". Here we have a variety of parameters, $C_1(t), C_2(t)$ for example, being set with constant proportion to certain statistics from the past — these constants are which the model could have learnt, though the search of which would take high computational power. An interesting extension would be to learn these parameters and discuss the improvement on the learning.

We also note the financial applicability of such an adaptive learning model: extra parameters can also be introduced to engage with practical application to trading, e.g. the loss function in Equation 12 could take a specification that relates to a rolling-averaged profit and loss, or a more realistic profit and loss with trading barriers, e.g. $\sum_{t=t_s+1}^{t_s+17} \text{sign}(\alpha_t) \mathbb{1}[\alpha_t \in R(\Phi_t, H)] \frac{P_{t+1} - P_t}{P_t}$ where $R(\Phi_t, H)$ is a time-varying region for the signal to be strong enough to trade, which can be one of the parameters being learnt. Such an extension may also engage with the contemporary econometric methods in conditional heteroskedasticity (Harvey 2013).

An additional direction is to engage with the study on penalisation, in particular the penalisation on estimation. In time series, due to the moving average terms, MLE is mostly inevitable, thus the penalisation must be done in the fashion of penalised MLE (pMLE). Here we focus on the penalisation on functional forms post-estimation while using the MLE at the first stage. An extension would be to engage with the theory of pMLE, e.g. from Cole, Chu, and Greenland (2013) and Spokoiny (2018), and use the adaptive learning to appreciate the value of penalisation in the context of forecasting. A particular modelling issue, as has been shown here, is the distaste towards multivariate models — this can be due to the failure of capturing the underlying time-varying parameters together with low degrees of freedom. This has also been studied by Wilms et al. (2017) and more thoughts on

penalisation and potentially "smart identification" using the past information could be worked on.

In conclusion, we propose a forecast-centric learning model that aims to adapt to the past information in a time series context. Such a model requires inputs of different functional forms, as well as training methods to produce parametric estimation and forecasting — these are handled by the traditional ARIMA models and explanatory variables which are generated from the order book. The result of the learning model is comparable to the top models if the models were to be fixed, and can outperform the fixed models in relatively volatile and non-stationary market conditions. Additionally, stability is more ensured and the error-handling, functional penalisation, and potentially other criteria can be encoded as part of the model selection process. Since the process is intuitive and interpretable, many extensions and applications can be made, for which we have shown an application to statistical testing, in both a Bayesian and a frequentist context.

A Additional statistical notes

A.1 On the ADF test

Here we adapt the following procedure for ADF Test with a constant. Let $\{y_t\}_{t=T_0}^{T_1}$ be the dataset we would like to run the test on. Regress

$$y_t = \mu + \sum_{j=1}^p \phi_j y_{t-j} + \varepsilon_t, \quad \varepsilon_t \sim iidN(0, \sigma^2), \quad t \in \{T_0 + p, \dots, T_1\} \quad (21)$$

The specific test is

$$H_0 : \phi_1 = 1$$

$$H_1 : \phi_1 < 1$$

Then construct test statistics $t := \frac{\hat{\phi}_1 - 1}{sd(\hat{\phi}_1)}$ where $\hat{\phi}_1$ and $sd(\hat{\phi}_1)$ are the MLE estimates and standard deviation estimates for ϕ_1 respectively from estimating Equation 21. The asymptotic result suggests²¹ $t \stackrel{H_0}{\sim} DF_{T_1 - T_0 - p}$. Thus, at a picked significance level, the critical value t^c is picked from $DF_{T_1 - T_0 - p}$ distribution and conclude.²²

- Reject H_0 and conclude stationarity if $t < t^c$;
- Do not reject H_0 and conclude non-stationarity if $t \geq t^c$.

There are a variety of choices of p one could take, usually $p \leq (T_1 - T_0 + 1)^{\frac{1}{3}}$ (Said and Dickey 1984). Here, to ensure consistency across different windows, we decide to put $p = 2$, as otherwise the small-window tests and later models get over-fitted.

A.2 On MAE and MSE

Let $\{y_{t+1|t}\}_{t \in T}$ be the forecasts and let $\{y_{t+1}\}_{t \in T}$ be the validation data. Write $card(T)$ as the cardinality of the index set T , then the MAE and MSE are:

$$MAE = (card(T))^{-1} \sum_{t \in T} |y_{t+1|t} - y_{t+1}|$$

$$MSE = (card(T))^{-1} \sum_{t \in T} |y_{t+1|t} - y_{t+1}|^2$$

²¹ See Fuller (1996).

²² Note: t^c is negative here.

B Additional figures and tables

	OFI mean	OFI p-score	OIBmean	OIB p-score	VWM	Number per bracket
count	5040	5040	5040	5040	5040	5040
mean	0.17	0.54	0.00	0.51	4068.21	424.79
std	0.10	0.02	0.06	0.06	141.09	84.88
min	-0.25	0.45	-0.39	0.20	3734.99	81
25% quantile	0.10	0.52	-0.03	0.47	3996.03	377
50% quantile	0.16	0.54	0.00	0.51	4052.03	431
75% quantile	0.23	0.55	0.04	0.55	4137.43	483
max	0.95	0.62	0.33	0.77	4421.84	601

Table 11: Summary of statistics for features (column 1 to 4), the VWM (column 5), and the number of observations per 5-minute bracket (column 6).

Model Group	0	1,3,7,9	2,3,8,9	4,6,10,12	5,6,11,12
Explanatory Variables	None	OIBmean	OFI mean	OIBp-score	OFI p-score

Table 12: Explanatory variables used in each of the fixed model groups

λ	1	0.99	0.95	0.9	0.85	0.8
$\log_{\lambda}(0.5)$	∞	68.97	13.51	6.58	4.27	3.11
λ^{47}	1	0.62	0.09	0.01	0	0

Table 13: Logarithmic and exponential properties amongst the choices of λ (Up to two decimal places)

Figure 18: Upper: histogram of the number of observations per bracket. Centre and Lower: histogram of the day gap, lunch gap, and the others, laid down on the same scale.

Figure 19: Bayesian testing result (left) and frequentist testing result (right)

References

- Aghabozorgi, Saeed, Ali Seyed Shirkhorshidi, and Teh Ying Wah (2015). “Time-series clustering—A decade review.” In: *Information Systems* 53, pp. 16–38.
- Akaike, H. (1974). “A new look at the statistical model identification.” In: *IEEE Transactions on Automatic Control* 19.6, pp. 716–723.
- Andres, Philipp and Andrew C Harvey (2012). “The Dyanamic Location/Scale Model: with applications to intra-day financial data.” In: *Cambridge Working Papers in Economics* 1240.
- Avellaneda, Marco, Josh Reed, and Sasha Stoikov (2011). “Forecasting prices from level-I quotes in the presence of hidden liquidity.” In: *Algorithmic Finance* 1, pp. 35–43.
- Banerjee, Anindya et al. (1993). *Co-integration, Error Correction, and the Econometric Analysis of Non-Stationary Data*. 1st ed. Oxford University Press.
- Cai, Zongwu and Xian Wang (2014). “Selection of Mixed Copula Model via Penalized Likelihood.” In: *Journal of the American Statistical Association* 109.506, pp. 788–801.
- Casella, George and Roger L Berger (2008). *Statistical Inference*. 2nd ed. Cengage Learning.
- Chakraborty, Chiranjit and Andreas Joseph (2017). “Machine learning at central banks.” In: *Cambridge Working Papers in Economics* 674.
- Cole, S. R., H. Chu, and S. Greenland (2013). “Maximum Likelihood, Profile Likelihood, and Penalized Likelihood: A Primer.” In: *American Journal of Epidemiology* 179.2, pp. 252–260.
- Cont, Rama, Arseniy Kukanov, and Sasha Stoikov (2014). “The Price Impact of Order Book Events.” In: *Journal of Financial Econometrics* 12.1, pp. 47–88.
- Cont, Rama, Sasha Stoikov, and Rishi Talreja (2010). “A Stochastic Model for Order Book Dynamics.” In: *Operations Research* 58.3, pp. 549–563.
- Fuller, Wayne A (1996). *Introduction to Statistical Time Series*. 2nd ed. Wiley.
- Goodfellow, Ian, Yoshua Bengio, and Aaron Courville (2016). *Deep Learning*. The MIT Press.
- Harris, Richard and Robert Sollis (2003). *Applied Time Series Modelling and Forecasting*. 1st ed. Wiley.
- Harvey, Andrew C (1989). *Forecasting, structural time series models and the Kalman Filter*. 1st ed. Reprint 2001. Cambridge University Press.
- (1993). *Time Series Models*. 2nd ed. Harvester Wheatasheaf.
- (2013). *Dynamic Models for Volatility and Heavy Tails: With Applications to Financial and Economic Time Series*. 1st ed. Econometric Society Monographs. Cambridge University Press.
- Harvey, Andrew C and Genaro Sucarrat (2014). “EGARCH models with fat tails, skewness and leverage.” In: *Computational Statistics & Data Analysis* 76, pp. 320–338.
- Hastie, Trevor, Robert Tibshirani, and Jerome Friedman (2001). *The elements of statistical learning: data mining, inference, and prediction*. 1st ed. Springer.
- Hochreiter, Sepp and Jürgen Schmidhuber (1997). “Long Short-Term Memory.” In: *Neural Computation* 9.8, pp. 1735–1780.
- Klenke, Achim (2013). *Probability Theory: A Comprehensive Course*. 2nd ed. Springer.

- Koop, Gary, Dale J. Poirier, and Justin L. Tobias (2007). *Bayesian Econometric Methods*. Cambridge University Press.
- Li, Zhenwei, Jing Han, and Yuping Song (2020). “On the forecasting of high-frequency financial time series based on ARIMA model improved by deep learning.” In: *Journal of Forecasting*. DOI: [10.1002/for.2677](https://doi.org/10.1002/for.2677).
- Prado, Raquel and Mike West (2010). *Time Series: Modeling, Computation, and Inference*. 1st ed. Taylor and Francis Group.
- Rubio, Gines et al. (2007). “Kernel Methods Applied to Time Series Forecasting.” In: *Computational and Ambient Intelligence*. Ed. by Francisco Sandoval et al. 1st ed. Springer, pp. 782–789.
- Said, Said E. and David A. Dickey (1984). “Testing for Unit Roots in Autoregressive-Moving Average Models of Unknown Order.” In: *Biometrika* 71.3, pp. 599–607.
- Shalev-Shwartz, Shai and Shai Ben-David (2014). *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press.
- Sirignano, Justin and Rama Cont (2018). *Universal features of price formation in financial markets: perspectives from Deep Learning*. arXiv:1803.06917.
- Spokoiny, Vladimir (2018). *Penalized maximum likelihood estimation and effective dimension*. arXiv:1205.0498v4.
- Stoikov, Sasha (2014). “Time is Money: Estimating the Cost of Latency in Trading.” In: *Market Microstructure: Confronting Many Viewpoints, Paris*.
- Tibshirani, Robert et al. (2005). “Sparsity and smoothness via the fused lasso.” In: *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 67.1, pp. 91–108.
- Vapnik, Vladimir N. (2000). *The Nature of Statistical Learning Theory*. 2nd. Springer.
- Wilms, Ines et al. (2017). *Sparse Identification and Estimation of Large-Scale Vector AutoRegressive Moving Averages*. arXiv:1707.09208.
- Yang, Parley Ruogu (2020). “Using The Yield Curve To Forecast Economic Growth.” In: *Journal of Forecasting*. DOI: [10.1002/for.2676](https://doi.org/10.1002/for.2676).
- Zbonakova, Lenka, Xinjue Li, and Wolfgang K. Härdle (2018). “Penalized Adaptive Forecasting With Large Information Sets and Structural Changes.” In: *SSRN Electronic Journal*. IRTG 1792 Discussion Paper 2018-039.