BANK OF GREECE
EUROSYSTEM

Working Paper

Trading VIX on volatility forecasts: Another volatility puzzle?

Stavros Degiannakis
Panagiotis Delis
George Filis
George Giannopoulos

336

FEBRUARY 2025

BANK OF GREECE
Economic Analysis and Research Department – Special Studies Division
21, E. Venizelos Avenue
GR-102 50 Athens
Tel: +30210-320 3610
Fax: +30210-320 2432

www.bankofgreece.gr

Published by the Bank of Greece, Athens, Greece

ISSN: 2654-1912 (online)
DOI: https://doi.org/10.52903/wp2025336

TRADING VIX ON VOLATILITY FORECASTS: ANOTHER VOLATILITY PUZZLE?

Stavros Degiannakis
Bank of Greece and Panteion University of Social and Political Science

Panagiotis Delis
Bank of Greece and University of Piraeus

George Filis
University of Patras

George Giannopoulos
Kingston University and University of West Attica

ABSTRACT

This study evaluates the economic usefulness of stock market implied volatility forecasts, based on their ability to improve the short-run trading decision-making process. The current literature aligns the forecast horizon with the frequency of the trading decision in order to evaluate different forecasting frameworks. By contrast, the premise of our paper is that these should not be necessarily related, but rather the evaluation should be based on the actual needs of the end-user. Thus, we evaluate whether the multiple days ahead stock market volatility forecasts vis-à-vis the 1-day ahead forecasts can improve the 1-day ahead trading profits from VIX and the S&P500 futures. Our results suggest that indeed the 1-day ahead trading profits are significantly improved when the trading decisions are based on longer-term volatility forecasts. More specifically, the highest trading gains are obtained when using the 22-days-ahead forecasts. The results hold true for both VIX and S&P500 futures day-ahead trading. Although there is no theoretical background regarding the fact that forecasting and trading horizons should not be aligned, we strongly motivate this potential issue, both from the statistical and financial point of views.

Keywords: Stock market implied volatility, volatility forecasts, trading profit, HAR model.

JEL classification: C22, C53, G11, G15, G17.

Acknowledgements: The views expressed in this paper are those of the authors and not necessarily those of either the Bank of Greece or the Eurosystem.

Correspondence:
Stavros Degiannakis
Bank of Greece,
21 E. Venizelos Avenue, 10250
Athens, Greece
email: sdegiannakis@bankofgreece.gr

1. Introduction

A growing body of literature recognizes the importance of stock market volatility forecasting. With respect to forecasting models, the existing research suggests that the Heterogeneous Autoregressive Model (HAR), proposed by Corsi (2009), is the most efficient model (i.e., Bollerslev and Wright, 2001; Andersen *et al.*, 2007; Busch *et al.*, 2011; Fernandes *et al.*, 2014; Taylor, 2014; Patton and Sheppard, 2015; Kourtis *et al.*, 2016; Liang *et al.*, 2020a; Kambouroudis *et al.*, 2021; Ye *et al.*, 2022). Nevertheless, for the evaluation of volatility forecasts, the research to date tends to focus on statistical (e.g., mean squared errors, quasi-likelihood, or success ratio) rather than on the economic purpose of the stock market volatility forecasts and more specifically for trading purposes.

In the volatility forecasting literature, the objective-based evaluation is generally understood as assessing the forecasts' economic usefulness (i.e., the purpose of the stock market volatility forecasts). As noted by Elliott and Timmermann (2008), relating the forecasts with the agents' economic decisions enhances the evaluation of the forecast for the end-user. Conversely, by considering the economic value of the investment, the forecast evaluation based on an objective-based criterion (i.e., economic criterion) allows the investors to understand the forecasting accuracy through the economic measurement units (e.g., trading profit) that are relevant to them (Taylor, 2014).

To date, studies use different objective-based criteria to evaluate the forecasts. For instance, Engle *et al.* (1993), Noh *et al.* (1994), Elder and Gannon (1998), Angelidis and Degiannakis (2008) and more recent studies (e.g., Andrada-Félix *et al.*, 2016; Degiannakis and Filis, 2017; Degiannakis *et al.*, 2018; Degiannakis and Filis, 2022; Delis *et al*., 2023) have assessed the volatility forecast using the straddles option trading strategy. On the other hand, Chou and Liu (2010), Kourtis *et al.* (2016) and Branco *et al.* (2024) have used a volatility-timing strategy, while optimal portfolio allocation context using the mean–variance framework is considered by Becker *et al.* (2015), Liang *et al*. (2020b), Ye *et al.* (2022), Ma *et al.* (2023), Son *et al.* (2023), Salisu *et al.* (2023) and Zhang *et al.* (2024). Furthermore, Jiang *et al*. (2014) and Chkili *et al.* (2014) examine the value-at-risk as an alternative objective-based criterion and Qiao *et al.*

(2022) use a simple VIX trading strategy. In addition, Engle *et al.* (1996) have used the ability to maximize realized trading profit to assess the volatility forecasts' quality.

Regarding the economic usefulness of stock market volatility forecasts, previous research suggests that realized volatility forecasts are less informative in turbulent periods when these are used for straddle trading (Andrada-Félix *et al.*, 2016). Furthermore, Engle *et al.* (1996) demonstrated a significant difference in trading profits between volatility forecasts obtained by various models. Nevertheless, Bernales and Guidolin (2014) analysis of the dynamics of the implied volatility surface suggests that trading profits diminish after the inclusion of transaction costs.

In terms of the implied volatility indices forecasts, recent evidence indicates that they have predictive power on the synchronization of stock market returns (Magner *et al.*, 2021). In particular, it has been suggested that Chicago Board Options Exchange Volatility Index[1] (CBOE VIX), European STOXX 50 Volatility (VSTOXX) and volatility Index Japan (VXJ) have high predictive power in the synchronization of stock market returns, though VIX prevails over the other indices (Magner *et al.*, 2021). In addition, Liang *et al.* (2020b) and Branco *et al.* (2024) demonstrated that investors could realize higher utility gains by using implied volatility models. Similarly, Kourtis *et al.* (2016) findings suggest the superiority of implied volatility over historical volatility methods for the diversification of the international portfolio.

In summary, it is rather important to evaluate forecasts not only based on statistical-loss functions but also based on their ability to improve a decision-making process, such as a trading decision. Although the literature has shown that research has been carried out on the economic usefulness of stock market volatility forecasts, there is still an important gap in the literature that the present study tries to fill. Typically, the existing literature matches the forecast horizon with the investment or trading horizon of the end-user of the volatility forecast. So, for instance, existing research assesses the *s*-days ahead volatility forecast using the outcome of the *s*-days ahead economic-based criterion (e.g., profits or losses in case where a trading strategy has been implemented). Nevertheless, we posit that the forecasting horizon should not be necessarily related to the investment horizon of the volatility forecast end-user. As far as our knowledge is concerned this is the first time that such a rule is applied. In particular, in this study the

[1] The VIX is the 30-day real-time volatility index of the S&P500.

forecasting period of the stock market implied volatility is a number of s-days ahead. However, the assessment of performance (trading profits) is based only on the 1-day ahead trading profit. To do so, we first estimate a HAR model for VIX (HAR-IV), which we further augment with the predictive information from various other implied volatility indices (HAR-IV-X). Then, we evaluate the performance of these implied volatility forecasts based on the 1-day ahead after-cost trading profits from two real world trading assets, namely, the VIX futures and the S&P500 futures.

In greater detail, in this study we produce stock market implied volatility forecasts for up to 66-days ahead, which we evaluate for their effectiveness in generating 1-day trading profits. Therefore, we assume that, under each model framework, there are 66 traders who are only interested in the 1-day ahead trading decisions. Trader no.1 generates 1-day ahead forecasts for VIX, using HAR-IV and HAR-IV-X models, and subsequently she will make her trading decision for VIX futures and S&P500 futures for the next day (trading day $t+1$). Equivalently, traders no.2, …, no.66 would generate 2-days, …, 66-days ahead forecasts for VIX, respectively and they will then proceed with their trading decision for the next day (trading day $t+1$).

Before we proceed any further, it would be important to clarify the rationale behind the application of this trading rule for the evaluation of stock market implied volatility forecasts.

We are aware that for any arbitrary trading day the forecast user computes the predictions for the out-of-sample period of s-trading days ahead; i.e., the days $t+1,\dots,t+s$, based on an in-sample period of size $\tilde{T}$; i.e. the days $t-\tilde{T}+1,\dots,t-1,t$. Let us also define as a trading period the time interval on which an investor wants to trade. So, for any arbitrary trading day, the investor trades for any s-trading days ahead; i.e., the days $t+1,\dots,t+s$.

Until now, the only approach used in financial and econometric literature is the matching between the forecasting and trading period. However, this practice has some drawbacks for a trading period of $s>1$, from both econometric and financial points of view.

From an econometric point of view, we have to deal with overlapping returns. For an out-of-sample period of s-trading days, every trading day is considered s times. Let us assume that the econometrician provides volatility predictions for s days ahead and the daily returns for the volatility trading are denoted as $y_{t+1}, .., y_{t+s}$. Each daily return is forecasted s times, thus, we have the conditional forecasts $y_{t+i\backslash t+i-j}$, for $i = 1, \dots, s$ and $j = 1, \dots, s$, given that $j \leq i$. For example, for $s = 3$ and the trading days named as $t + 1$ =Monday, $t + 2$ =Tuesday, and $t + 3$ =Wednesday, we have 3 forecasts for each day. Indicatively, for the $t + 3$ = Wednesday, we compute the 1-day ahead forecast $y_{t+3\backslash t+2}$, the 2-day ahead forecast $y_{t+3\backslash t+2}$ and the 3-day ahead forecast $y_{t+3\backslash t}$. Henceforth, the rolling estimations are based on s overlapping predictions. The overlapping observations assume different statistical inference for the estimation of the standard errors of the time series under investigation.

The asymptotic inference for overlapping returns has been investigated by Hansen and Hodrick (1980) and Richardson and Stock (1989). Engle *et al.* (1996) computed the overlapping rate of returns of a trading strategy based on options and showed that the Hansen-Hodrick standard error is ten times greater than the ordinary standard error. The overlapping observations induce strong bias (i.e., for $s = 10$, if the average rate of return for the whole period is around 1.5% but there is an outlier rate of return equal to 8%, then this outlier will affect $s = 10$ predictions in total).

From a financial point of view, the investor uses an initial amount of capital, defined as k, in order to fulfil the trading orders imposed by the forecasting signals. Let us assume that $k = \$1{,}000$. For a trading period of $s = 1$, the trader opens a trading position of \$1,000 the first trading day and closes this position the second trading day. On the second trading day, she opens a new position using the capital available on that day, $\$1{,}000(1 + y_{t+1})$ (i.e., plus/minus the profits/losses of the previous position). Subsequently, the third trading day, she closes the position of the second trading day and opens a new position of $\$1{,}000(1 + y_{t+1})(1 + y_{t+2})$, and so on. Let us assume now that we have a trading period of $s = 3$ and the same initial amount of capital $k = \$1{,}000$. Now, the trader opens a trading position of $\$1{,}000/s$ the first trading day. The second and third trading days she must also open a trading position of $\$1{,}000/s$. Then the forth trading day, the trader must close the position opened the first day and at the same time she must open a position for that day of $\$1{,}000(1 + y_{t+1})/s$. It is not rational at all, for an investor to handle at the same time so many multiple open positions neither from a rational expectation nor from a risk management point of view. Having multiple open positions equals to multiplying the exposure to risk. If, for example, there is a trading day with extreme but plausible negative returns, then this loss will follow the trader for s times.

Another approach that has been applied in the literature is the computation of the trading profits for non-overlapping points in time. This approach has the advantage that we avoid the biasness induced by the autocorrelation of the forecast errors across time, but it also has the disadvantage that the investor is not able to trade at each trading day, which is too restrictive. For example, the investor would be able to trade only a fraction $s/\breve{T}$ of the total trading period, where $\breve{T}$ denotes the total number of out-of-sample trading days.

Hence, a more real-world approach, which would be more applicable to market participants is to identify whether the 1-day ahead volatility forecasts can indeed generate superior 1-day ahead trading profits, or whether the s-days ahead volatility forecasts could provide better results. A secondary but also important aspect of our study is the examination of whether there are any specific predictors that are capable of producing even higher trading profits.

The findings of this study provide material evidence in favor of our novel hypothesis. In particular, the 1-day ahead trading gains are significantly higher when such trading decisions are based on the 22-days ahead stock market implied volatility forecasts, as opposed to the 1-day ahead forecasts. Moreover, another finding that could attract the attention of the professional forecasters and traders is that the above-mentioned highest trading gain is obtained using the HAR-IV-TYVIX (10-year US Treasury note Volatility Index) model in terms of both VIX and S&P500 futures trading. The fact that the 10-year US Treasury note Volatility Index representing the macroeconomic conditions can help improving the VIX and S&P500 futures trading outcome can be considered vital for forecasters and traders.

Although there is no theoretical background regarding the fact that forecasting and trading horizons should not be the same, we strongly motivate this potential issue, both from the statistical and financial point of views. Our findings, based on the framework that the paper is developing, highlight the necessity for such theory to be developed.

The rest of the paper is structured as follows. Section 2 presents the data used in the study. Section 3 discusses the modelling framework. The adopted forecasting procedures are provided in Section 4, whereas Section 5 analyses the findings. Finally, the paper is concluded in Section 6.

2. Data description

As already mentioned in Section 1, in this paper we generate VIX forecasts and use them on trading VIX and S&P500 index futures. VIX historical data have been retrieved by CBOE[2]. The VIX and the S&P500 index futures have been obtained by CBOE/CFE and Bloomberg.

We should mention that VIX forecasts are generated using models that include information not only from VIX itself but also from implied volatility indices that represent other asset classes and markets. The scope of using these indices as predictors of VIX is to enhance the forecasting performance of the benchmark models used for generating VIX forecasts.

Thus, the relevant predictors that we use consist of the British Pound Volatility Index (BPVIX), the Euro Volatility Index (EUVIX), the Yen Volatility Index (JYVIX), the Crude oil ETF volatility index (OVX), the 10-year US Treasury note Volatility Index (TYVIX), the DAX Volatility Index (VDAX), the Nikkei Volatility Index (VNIK) and the EURO STOXX 50 Volatility Index (VSTOXX). The source, which is used for retrieving historical data is the Investing[3].

The sample period spans from the 10th of May, 2007 up to the 3rd of April, 2020 (3250 trading days) and the frequency of the dataset is daily. The in-sample period runs from 10th of May, 2007 to 26th of April, 2011 while the remaining days are used for our out-of-sample period. During the out-of-sample period, the models are re-estimated at each trading day, having in mind to avoid the look-ahead bias; i.e. the use of information that would not have been available during the period being analyzed.

However, we note that for our additional forecast evaluation strategy, the sample period is extended up to 18th of February, 2022 (475 more trading days). In more detail, we use the period between the 6th April 2020 and 18th February 2022, as our post out-of-sample period, in order to investigate whether the forecasting ability remains

[2] The calculation of VIX is based on the CBOE methodology. See https://www.cboe.com for further details.
[3] The data that support the findings of this study are available from the corresponding author upon request.

qualitatively the same. In forecasting literature, an important issue that often remains unanswered is the problem of data snooping. Data snooping, or data fishing, refers to testing hypotheses once you have seen the data. In our case of forecasting evaluation, data fishing is the phenomenon of finding a model that produces accurate forecasts only for the out-of-sample period under investigation. In the present study, we decided to evaluate the models that have the best predictive ability in a sample in which they have not been evaluated before. Hence, if those models retain the same forecasting performance, then we can refer that they do not suffer from the data snooping.

Table 1 reports the descriptive statistics of the implied volatility indices. It can be easily observed that VIX presents the highest value of coefficient of variation (CV), followed by OVX and VSTOXX, suggesting that these indices have a much greater dispersion relatively to other implied volatility indices. Furthermore, we note that all indices are positively skewed and exhibit excess kurtosis, as expected. Moreover, from Figure 1, we observe that the values of all implied volatility indices are increased during the Global Financial Crisis of 2007-2009, which shows the uncertainty that existed during that period in all asset markets. Finally, we notice that OVX assumes its higher values (almost 200) during the COVID-19 pandemic, which can be explained by the huge interruptions in the global oil market during that period.

[TABLE 1 HERE]

[FIGURE 1 HERE]

3. Modelling framework

Simple Heterogeneous Autoregressive model (HAR-IV)

As far as we are concerned, the HAR model specification proposed by Corsi (2009) is the most widely used methodology for generating realized and implied volatility forecasts (e.g., Fernandes *et al.*, 2014). In this study, we first use the simple HAR model specification without including any exogenous information. Therefore, in case of VIX, the HAR-IV model is written as follows:

$$log(VIX_t) = a_{0,t} + a_{1,t} log\left(VIX_{t-1}^{(d)}\right) + a_{2,t} log\left(VIX_{t-1}^{(w)}\right) + a_{3,t} log\left(VIX_{t-1}^{(m)}\right) + \varepsilon_t, \tag{1}$$

where $a_{0,t}$, $a_{1,t}$, $a_{2,t}$ and $a_{3,t}$ are the coefficients to be estimated and ε_t is the error term. Regarding the lag terms of the simple HAR model specification, $VIX_{t-1}^{(d)}$ represents the first lag of VIX, $VIX_{t-1}^{(w)}$ is the weekly component, which is computed as $5^{-1}\sum_{k=1}^{5} log(VIX_{t-k})$ and $VIX_{t-1}^{(m)}$ represents the monthly component, which is computed as $22^{-1}\sum_{k=1}^{22} log(VIX_{t-k})$.[4]

Heterogeneous Autoregressive model with exogenous predictors (HAR-IV-X)

In addition to the benchmark HAR-IV model, we estimate a HAR-IV-X model including each of the implied volatility measures as exogenous predictors. For example, the HAR-IV-OVX model includes apart from the HAR lag components an additional explanatory variable, which is the first lag of OVX. In general, the model entitled HAR-IV-X is estimated in the form:

$$\begin{aligned} log(VIX_t) = a_{0,t} + a_{1,t} log\left(VIX_{t-1}^{(d)}\right) + a_{2,t} log\left(VIX_{t-1}^{(w)}\right) \\ + a_{3,t} log\left(VIX_{t-1}^{(m)}\right) + a_{4,t} log\left(X_{(1),t-1}^{(d)}\right) + a_{5,t}\left(X_{(1),t-1}^{(w)}\right) \\ + a_{6,t}\left(X_{(1),t-1}^{(m)}\right) + \varepsilon_t, \end{aligned} \qquad (2)$$

where $X_{(1),t}$ expresses the various implied volatility measures; i.e.: $X_{(1),t}$: $\{BPVIX_t, EUVIX_t, JYVIX_t, OVX_t, TYVIX_t, VDAX_t, VNIK_t, VSTOXX_t\}$. More specifically, $X_{(1),t-1}^{(d)}$ represents the first lag of $X_{(1)}$, $X_{(1),t-1}^{(w)}$ is the weekly component of $X_{(1)}$, which is computed as $5^{-1}\sum_{k=1}^{5} log\left(X_{(1),t-k}\right)$, and $X_{(1),t-1}^{(m)}$ represents the monthly component of $X_{(1)}$, which is computed as $22^{-1}\sum_{k=1}^{22} log\left(X_{(1),t-k}\right)$.

4. Forecasting framework

The computation of the predictions should be based only on the information set that is available up to the current trading day *t*. Otherwise, the look ahead bias phenomenon would have appeared. The forecasting procedure should not be a black

[4] As Bollerslev *et al.* (2016) proposed, the asymptotic theory for high-frequency volatility estimators suggests the incorporation of realized quarticity in the HAR model. In the case that we had modelled realized volatility instead of implied volatility indices, the HAR model should have been tested against the HARQ model, to investigate whether the standardization of volatility with the quarticity leads to more accurate predictions.

box. Thus, we provide a representative example of a trader should compute the s-trading days ahead forecast of the HAR-IV-X models:

$$VIX_{t+s|t} = exp\Big(\hat{a}_0^{(t)} + \hat{a}_1^{(t)} log\big(VIX_{t+s-1|t}\big) + a_2^{(t)}\big(s^{-1}\sum_{k=1}^{s-1} log\big(VIX_{t-k+s|t}\big) + (5-s)^{-1}\sum_{k=s}^{5} log(VIX_{t-k+s})\big) + \hat{a}_3^{(t)}\big(s^{-1}\sum_{k=1}^{s-1} log\big(VIX_{t-k+s|t}\big) + (22-s)^{-1}\sum_{k=s}^{22} log(VIX_{t-k+s})\big) + \hat{a}_4^{(t)} log\big(X_{(1),t+s-1|t}\big) + \hat{a}_5^{(t)}\big(s^{-1}\sum_{k=1}^{s-1} log\big(X_{(1),t-k+s|t}\big) + (5-s)^{-1}\sum_{k=s}^{5} log\big(X_{(1),t-k+s}\big)\big) + \hat{a}_6^{(t)}\big(s^{-1}\sum_{k=1}^{s-1} log\big(X_{(1),t-k+s|t}\big) + (22-s)^{-1}\sum_{k=s}^{22} log\big(X_{(1),t-k+s}\big)\big) + {}^{1}\!/_{2}\,\hat{\sigma}_{\varepsilon}^{2}\Big). \quad (3)$$

The $VIX_{t-k+s|t}$ is the prediction of VIX for trading day $t-k+s$ based on the information set that is available up to trading day t. Similarly, the $X_{(1),t-k+s|t}$ terms represent the prediction for the implied volatility index $X_{(1)}$ for trading day $t-k+s$ based on the information set that is available up to trading day t. The $X_{(1),t-k+s}$ is the actual price of the implied volatility index $X_{(1)}$ on the trading day $t-k+s$. Finally, the conditional forecasts of the $X_{(1)}$ are being computed by satellite HAR models.

4.1. Forecasting strategy

We use an initial sample period of $\tilde{T} = 1000$ trading days. The rest $\breve{T} = 2250$ trading days are used for the real out-of-sample forecasting period. The choice of the initial sample period is justified by the fact that a large enough sample size is required for the estimation of the forecasting models but also due to the fact we intentionally need the post-2014 period to be part of the out-of-sample period[5]. We produce forecasts from 1-day up to 66-days ahead. Hence, for the first set of real-out-of-sample forecasts for 1-day to 66-days ahead, we use the initial sample period $\tilde{T} = 1000$. For the remaining forecasts we employ a rolling window approach with a fixed window length of 1000 daily observations.

[5] We have also used 750 and 1250 trading days for our initial sample period for robustness purposes and the results remain qualitatively similar. The results are available in the Appendix.

4.2. Forecast evaluation criteria

We proceed to the evaluation of the forecasting performance based on economic loss functions. A similar approach was employed by Delis *et al.* (2022), who's forecasting evaluation technique was based on a quasi-trading strategy.

<u>*Trading strategy 1 (TS1)*</u>:

If $VIX_{(j)t+s|t} > VIX_t$, for $s = 1, \dots, 66$ trading days ahead, then the trader, who follows the (j) model predictions, takes long position on futures of VIX index at trading day t, VIX_t^F. Of course, if $VIX_{(j)t+s|t} \leq VIX_t$ then the trader takes short position on futures of VIX index at trading day t. For each one of the $j = 1, \dots, 9$ forecasting models the cumulative returns (CR) from TS1 are computed as:

$$CR_{(j)}^{(s)} = \sum_{t=1}^{\tilde{T}} \left(I_{(j),t} \times \frac{(VIX_{t+1}^F - VIX_t^F)}{VIX_t^F} \right), \tag{4}$$

where $I_{(j),t} = \begin{cases} 1 & \text{if} \quad VIX_{(j)t+s|t} > VIX_t \\ -1 & \text{if} \quad VIX_{(j)t+s|t} \leq VIX_t \end{cases}$ and (s) denotes the trading days ahead.

<u>*Trading strategy 2 (TS2)*</u>: Trade the underlying price of the S&P500 (as well as the S&P500 futures) based on implied volatility forecasts. The implied volatility index is considered by the market participants as the fear index. Hence, when investors expect volatility to increase (decrease) then they anticipate that the underlying stock index will decline (increase). Under this premise, if $VIX_{(j)t+s|t} > VIX_t$, for $s = 1, \dots, 66$ trading days ahead, then the trader who follows the (j) model predictions takes short position on futures of S&P500 index at trading day t. Naturally, if $VIX_{(j)t+s|t} \leq VIX_t$ then the trader takes long position on futures of S&P500 index at trading day t. For each one of the j forecasting models the cumulative returns (CR) from TS2 are computed as:

$$CR_{(j)}^{(s)} = \sum_{t=1}^{\tilde{T}} \left(I_{(j),t} \times \frac{(S\&P500_{t+1} - S\&P500_t)}{S\&P500_t} \right), \tag{5}$$

where $I_{(j),t} = \begin{cases} 1 & \text{if} \quad VIX_{(j)t+s|t} \leq VIX_t \\ -1 & \text{if} \quad VIX_{(j)t+s|t} > VIX_t \end{cases}$ and (s) denotes the trading days ahead.

Once all aforementioned steps are followed, we evaluate the forecasting models based on the after-cost profits for the $t + 1$ trading day. We estimate the transaction costs per trade to be around 0.033%. As far as the VIX futures trading, the actual cost is \$6 per contract. We note that transaction costs are incurred every time a trader changes her position (from long to short and vice versa), as dictated by the VIX forecast for each of the s-days ahead horizons.

We should emphasize at this point that the trading must be conducted on the futures instead of the indices themselves, since the VIX index as well the S&P500 index are not tradeable assets. On the contrary the futures contracts on VIX and S&P500 are actual tradeable assets.

As a final step, we investigate whether there is statistically significant difference among the trading profits based on Hansen's *et al.* (2011) Model Confidence Set (MCS). The MCS identifies the set of the best models, where best is defined in terms of a predefined evaluation criterion, which is the trading profits in our case.

More specifically, the MCS explores the predictive ability of an initial set of M^0 models and investigates, at a predefined level of significance, which group of models survive an elimination algorithm. Let us define as $\Psi_{j,t}$ the evaluation function of model j at trading day t, and $d_{j,j^*,t} = \Psi_{j,t} - \Psi_{j^*,t}$ as the evaluation differential for $j, j^* \in M^0$. Given that the trading profits are under examination in our case, the evaluation function will be as follows: $\Psi_{j,t} = \left| I_{(j),t} \times \frac{(VIX_{t+1}^F - VIX_t^F)}{VIX_t^F} - \max_{(j)} \left(I_{(j),t} \times \frac{(VIX_{t+1}^F - VIX_t^F)}{VIX_t^F} \right) \right|$, for $I_{(j),t} = \begin{cases} 1 & \text{if} \quad VIX_{(j)t+s|t} > VIX_t \\ -1 & \text{if} \quad VIX_{(j)t+s|t} \leq VIX_t \end{cases}$. In short, when the MCS test is applied in the trading profits, the test shows whether there are additional models that can generate equal trading performance with the best performing model.

4.3. Risk-adjusted trading strategy performance

Apart from the trading profits that are used for the evaluation of the forecasting performance, we further implement a battery of risk-adjusted measures to assess the trading strategy performance. In this regard, we use well-established metrics that provide evidence of a well performed trading strategy.

The first metric that is computed for this purpose is the Sharpe ratio, which divides the portfolio's excess returns by a measure of its volatility[6]. In addition, we compute the Sortino ratio, which differs from the Sharpe ratio as it captures the standard

[6] It is noted that the portfolio of our study consists of only one asset, which is either S&P500 futures or VIX futures. Moreover, the effective Federal Funds rate is used for approximating the risk-free rate.

deviation of the downside risk rather than that of the total risk. We finally implement the Calmar ratio, which is measured as the portfolio's excess return over the maximum drawdown[7]. A high ratio implies that the returns of the trading strategy were not at risk of significant drawdowns. Contrarily, a low ratio indicates that the risk of drawdown is higher.

5. Empirical results

Before we proceed with the analysis of our findings, we shall reiterate here that in this paper we assume that there are 66 traders who are only making investment decisions for the next day. Hence, traders no.1, no. 2, up to no. 66 forecast VIX for 1-day, 2-days, …, 66-days ahead, respectively, using Eqs.1-2. Subsequently, each trader proceeds with her trading decision either for VIX futures (trading strategy 1 – TS1) or S&P500 futures (trading strategy 2 – TS2), given her VIX forecast.

Typically, the current literature, which mainly uses statistical-loss functions, tends to assess the statistical significance of the findings. By contrast, in our case the statistical significance among the different volatility forecasts is rather unimportant. This is so, as our focus is on trading profits. Real-world traders seek to identify the forecasting models that will maximise their trading profits. Hence, even if the best forecasting model is capable of generating a marginal profit (yet not significantly different from the second-best model), then a trader would not be indifferent between these two alternatives.

Finally, we note that the evaluation of the trading profits from each forecasting model should be on a risk-adjusted basis. To do so, we develop a quasi-risk evaluation of our strategies, which is related to the percentage of long and short positions that each trader takes for the VIX and S&P500 futures. Tables 2 and 3 reports these positions.

[TABLES 2 and 3 HERE]

As shown in Tables 2 and 3 traders who are using different forecasting horizons and different models tend to exhibit very similar percentage of long positions, suggesting that they are exposed in extremely similar risk level. Hence, the trading profits evaluation is based on equal risk terms.

[7] The maximum drawdown is an indicator that reflects the downside risk over a specific timeframe. It is calculated as the maximum observed loss from a peak to a trough of the portfolio.

5.1. Trading profits from VIX futures

We start our analysis with the trading profits from TS1. The results are shown in Table 4 and Figure 2.

[TABLE 4 HERE]

[FIGURE 2 HERE]

Table 4 shows indicatively the trading profits from VIX futures for the traders who produce forecasts for the next day (1-day ahead horizon), the following week (5-days ahead horizon), biweekly (10-days ahead horizon), month (22-days ahead horizon), two months (44-days ahead horizon) and three months (66-days ahead horizon). We note that the trader that focuses on the 22-days ahead VIX forecasts is capable of generating superior trading profits, relatively to all other horizons and models, as long as she incorporates in the forecasting model the predictive information from the 10yr US treasury notes implied volatility (HAR-IV-TYVIX). This finding is also strengthened by the MCS test that shows that the only model that survives the elimination algorithm is the HAR-IV-TYVIX model for the 22-days ahead forecasting horizon. Thus, there is no other model that has equal trading performance with the HAR-IV-TYVIX model.

The fact that the implied volatility of the US 10yr bonds provides superior trading profits can be explained by the close association between the two markets. In particular, TYVIX reflects uncertainty in the sovereign bond market, which further reflects potential illiquidity issues in the interbank market, credit risk, as well as, macroeconomic turbulence. These events are expected to result in heightened stock market uncertainty and as such it could operate as a leading indicator of stock market volatility (Welch and Goyal, 2008; Christiansen *et al.*, 2012; Gong *et al.*, 2023). Even more, investors utilise the bond market's information for their investment decisions. Hence, heightened volatility in the bond market forces stock market participants to require higher risk premium in the anticipation of higher economic risks, which would, in turn, lead to higher stock market volatility[8].

Another notable finding that we can observe from Table 4 is the fact that the trader that bases her trading decisions on the 1-day ahead VIX forecasts generate materially inferior profits, irrespectively the forecasting model, relatively to the traders

[8] The positive relationship between equity premium and stock market volatility has been established by several authors, including Bekaert and Wu (2000), Kim *et al.* (2004) and Gu *et al.* (2020), among others.

that produce longer-term forecasts. For instance, the higher profit level for the 1-day ahead horizon amounts to 790%, whereas the best performing model and forecast horizon can generate profits that amount to 988.94%. Even more, we show that the highest profit at any single forecasting model is achieved at a particular *s*-day ahead horizon rather than at the 1-day ahead, strengthening our main hypothesis that longer term forecasts are capable of providing superior information for the trading decisions of a day-ahead trader.

Figure 2 provides a more detailed view of the aforementioned findings. In short, we document that for next day trading on the VIX futures, a trader should proceed to longer-term VIX forecasts, rather than 1-day ahead. Furthermore, the use of exogenous predictive information from other implied volatility indices seems to provide an edge, since the higher profits from the HAR-IV model are 957.17% (at the 14-days ahead horizon[9]), whereas, as already mentioned, the highest profits are generated using the HAR-IV-TYVIX model for 22-days ahead horizon.

5.2. Trading profits from S&P500 futures

Next, we focus on the S&P500 futures trading strategy (TS2). The results are exhibited in Table 5 and Figure 3.

[TABLE 5 HERE]

[FIGURE 3 HERE]

As in the case of the VIX futures, Table 5 shows even more convincingly that the trader who makes her trading decisions based on the 1-day ahead horizon, irrespectively of the forecasting model, cannot achieve higher profits compared to the traders that use longer term forecasting horizons. In particular, the highest profits that this trader can obtain is 32.36%, using the HAR-IV-TYVIX model. However, the trader who focuses on the 22-days ahead forecasts, based on the HAR-IV-TYVIX model is capable of generating twice as much these profits, i.e., 66.18%. As in the case of the VIX futures trading, the HAR-IV-TYVIX model for the 22-days ahead forecasting

[9] In Table 2, contrary to Figure 2, the reader can see that the maximum profits for HAR-IV are 917.41% in the 22-days ahead horizon. This happens as in the tables we only show the standard *s*-days ahead horizons (i.e., 1-, 5-, 10-, 22-, 44- and 66-days ahead), whereas in the figures we show all *s*-days ahead from 1- up to 66-days ahead.

horizon is the only model that generates statistically superior profits against all other models and horizons, based on the MCS test.

Turning to Figure 3, we observe a forecast horizon that could yield even greater profits for a given trader, compared to the 66.18%. In particular, the trader who produces 6-days ahead VIX forecasts, so to make her trading decision for the following day on the S&P500 futures, can increase further her profits to the levels of 78.42%. Once again, the model that yields this result is the HAR-IV-TYVIX. Furthermore, as in the case of VIX futures trading, the highest profit at any single forecasting model is not achieved using the 1-day ahead forecasts but rather a particular *s*-day ahead prediction.

5.3. Forecast evaluation based on statistical-loss functions

In this section we report the findings of the forecasts' evaluation based on two statistical loss-functions, i.e. Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE), so to confirm that the latter functions can indeed lead to sub-optimal trading decisions vis-à-vis the economic-based loss functions. Tables 6 and 7 report the results from the RMSE and MAE, respectively, along with the MCS test[10].

[TABLES 6 and 7 HERE]

From Tables 6 and 7 we can notice that even though the HAR-IV-TYVIX is the best performing model for the 10-days up to 66-days ahead forecast horizons, it does not produce the lowest 1-day ahead forecasts. By contrast, the HAR-IV generated the lowest forecast error across all models and horizons. Hence, an investor that is interested in the 1-day ahead trading, either on VIX or the S&P500 futures, would base her trading decision based on the 1-day ahead forecasts of the HAR-IV model. However, as shown in Tables 4 and 5, her profits based on the HAR-IV model would be 790% and 17.75% for the VIX and S&P500 futures, respectively. These profit levels are materially lower compared to the ones that she would be able to generate if she had

[10] As mentioned in Section 4.2., the MCS explores the predictive ability of an initial set of models. For example, when the evaluation function is the squared forecast error, $\Psi_{j,t} \equiv \left(VIX_{(j)t+s|t} - VIX_{t+s}\right)^2$, then for $\vee\ j, j^* \in M$, $M \subset M^0$, the null hypothesis $H_{0,M}: E\left(d_{j,j^*,t}\right) = 0$ is tested against the alternative $H_{1,M}: E\left(d_{j,j^*,t}\right) \neq 0$, for some $j, j^* \in M$. The MCS test is also applied for the MAE statistical loss-function, as well.

based her trading decisions on 22-days ahead forecasts from the HAR-IV-TYVIX model (998.94% and 66.18% for VIX and S&P500 futures, respectively). Hence, the statistical-loss functions lead to sub-optimal trading decisions. As such, investors should evaluate their forecasts based on economic-based criteria that match the purpose of generating those forecasts.

5.4. Further assessment of the trading profits

Thus far we have established what might be potentially another volatility puzzle. Our results show for the first time a very important finding, namely, that day-ahead stock market traders should not base their trading decisions upon the next-day forecasts (1-day ahead horizon). On the contrary, producing longer-term forecasts tend to provide more valuable information for the traders as to which trading position they should assume at time t on the VIX or S&P500 futures. Even more, we show that augmenting the simple HAR-IV model to incorporate the predictive information of other assets' implied volatilities (HAR-IV-X), yield even greater profits, especially the information that is obtained from the 10yr US treasury notes implied volatility (TYVIX).

To elaborate further on these findings, Figure 4 shows the superior profits generated by the HAR-IV-TYVIX, over our out-of-sample period, compared to the HAR-IV for both the VIX futures and S&P500 futures trading.

[FIGURE 4 HERE]

It is evident that the trader who bases her decisions for the 1-day ahead trading on the 22-day ahead forecasts can materially improve her profitability compared to the trader that produces 1-day ahead forecasts in order to make her 1-day ahead trading decisions.

The above-mentioned evidence that longer-term forecasts could lead to higher trading performance than that of the 1-day ahead forecasts, is also valid when assessing the trading strategies (TS1 and TS2) performance based on a battery of risk-adjusted measures, as shown in Figure 5.

[FIGURE 5 HERE]

Figure 5 clearly shows that the 22-days ahead forecasts based on the HAR-IV model, which includes the information from the 10yr US treasury notes implied volatility (TYVIX) as predictor, generates superior risk-adjusted trading performance, based on the Sharpe, Sortino and Calmar ratios.

5.5. Post out-of-sample period

Overall, from Sections 5.1 and 5.2 we show VIX forecasting are indeed economically useful for traders that either invest their funds in the S&P500 futures or the VIX futures. Furthermore, we convincingly show that irrespectively of the forecasting model (HAR-IV or HAR-IV-X), in none of the cases the 1-day ahead forecasts are capable of generating higher profits compared to at least one *s*-days ahead forecasts. Thus, traders with one day trading horizon would achieve higher profits when their decisions are based on an *s*-day ahead volatility forecast (preferably the 22-days ahead) rather than the 1-day ahead.

However, even though we have obtained the aforementioned results based on the out-of-sample exercise, we further test whether our findings remain robust for a period that extends further out from our out-of-sample period. Put it simply, this section answers a rather important question for investors, namely, how could an investor make an informed *ex ante* choice of the best *s*-days ahead forecast. Thus, in this section we extend the comparison that was performed in Figure 4 for the two years following our out-of-sample period. More specifically, we assume that a trader has evaluated the outcomes of her models up to the present day (in this case this would be the last day of the out-of-sample period, i.e., 3rd April 2020). Subsequently, she chooses not to re-estimate her models but rather to follow the estimated best *s*-days ahead forecast, for the coming two years. Hence, we evaluate the day-ahead trading profitability based on the 22-days ahead forecasts relatively to the 1-day ahead forecasts, for the period 6th of April, 2020 up to 18th of February, 2022. As shown in Figure 6, there are significant gains even after the evaluation period (signaling the absence of data snooping) and even during a period of extreme events, since our post out-of-sample period covers the Covid-19 pandemic period, which caused significant turbulence in the financial markets.

[FIGURE 6 HERE]

Hence, the results obtained from this section strengthen our important findings that the *s*-days ahead forecasts are capable of generating incremental profits for the day-ahead trader, compared to the 1-day ahead forecasts. This is the first time that such a finding is reported in the related forecasting literature.

6. Conclusion

The aim of this study is to evaluate the economic usefulness of forecasted stock market implied volatility, based on their ability to improve a decision-making process, which in our study is a trading decision. Contrary to the common approach of previous studies that evaluate the *s*-days ahead forecasts based on their *s*-days ahead trading profits, we focus on the ability of the *s*-days ahead forecasts to generate 1-day ahead trading profits. In particular, our study evaluates the *s*-days ahead VIX forecasts on the level of the 1-day ahead trading profits from VIX and the S&P500 futures.

The rationale behind the proposition to generate 1-day ahead trading profits based on *s*-days ahead forecasts lies on the need to avoid (i) the biasness induced by the autocorrelation of the forecast errors due to the existence of overlapping returns (from an econometric point of view) and (ii) the existence of multiple open positions simultaneously (from a financial point of view).

Our results provide strong evidence that the forecasting horizon should not be related to the investment horizon of the volatility forecast end-user. More specifically, we show that a trader can maximise her 1-day ahead trading profits from the VIX and S&P500 futures when her trading decision are based on the 22-days ahead VIX forecasts. These results also hold true when applying risk-adjusted metrics as well as during the post out-of-sample experiment. Even more, we show that traders should forecast the 22-days ahead VIX by considering the predictive content of the US bond market implied volatility (TYVIX), which is capable of capturing the uncertainty level of the US economy.

Such findings have important implications for volatility trading. In particular, the results of this study directly affect the professional forecasters and traders, who assume synchronisation of forecasting and trading horizons, in the case of VIX futures trading. More specifically, we provide evidence that the forecasting and trading horizons should not be the same.

We acknowledge in the paper that there is no theoretical background regarding the fact that forecasting and trading horizons should not be the same. However, we strongly motivate this potential issue, both from the statistical and financial point of views. What is more important, is that our findings, based on the framework that the paper is developing, pave the way for future research to try to build the theory. It highlights the necessity for such theory to be developed. Let us not forget that many finance theories have emerged from pure empirical evidence, with most prominent being the EMH and CAPM, as well as, the behavioral finance. In addition, the ARCH model was also motivated by the initial empirical findings that Engle made in the early 1980s about volatility clustering. Finally, research community could be motivated by the outcome of this paper and further evaluate the proposed rationale by implementing this methodology in other asset classes and by using alternative trading strategies.

References

Andersen, T. G., Bollerslev, T. and Diebold, F.X. (2007). Roughing it up: Including jump components in the measurement, modeling and forecasting of return volatility, *The Review of Economics and Statistics*, 89(4), pp. 701-720.

Andrada-Félix, J., Fernández-Rodríguez, F. and Fuertes, A. (2016). Combining nearest neighbor predictions and model-based predictions of realized variance: Does it pay? *International Journal of Forecasting*, 32, pp. 695-715.

Angelidis, T. and Degiannakis, S. (2008). Volatility forecasting: Intra-day versus inter-day models, *Journal of International Financial Markets, Institutions and Money*, 18, pp .449-465.

Becker, R., Clements, A.E., Doolan, M.B. and Hurn, A.S. (2015). Selecting volatility forecasting models for portfolio allocation purposes, *International Journal of Forecasting*, 31(3), pp. 849-861.

Bekaert, G. and Wu, G. (2000). Asymmetric volatility and risk in equity markets, *The Review of Financial Studies*, *13*(1), 1-42.

Bernales, A. and Guidolin, M. (2014). Can we forecast the implied volatility surface dynamics of equity options? Predictability and economic value test, *Journal of Banking and Finance*, 46, pp. 326-342.

Bollerslev, T., Wright, J.H. (2001). High-frequency data, frequency domain inference, and volatility forecasting, *The Review of Economics and Statistics*, 83(4), pp. 596-602.

Bollerslev, T., Patton, A.J. and Quaedvlieg, R. (2016). Exploiting the errors: A simple approach for improved volatility forecasting, *Journal of Econometrics*, 192(1), 1-18.

Branco, R.R., Rubesam, A. and Zevallos, M. (2024). Forecasting realized volatility: Does anything beat linear models?, *Journal of Empirical Finance*, 101524.

Busch, T., Christensen, B.J. and Nielsen, M. Ø. (2011). The role of implied volatility in forecasting future realized volatility and jumps in foreign exchange, stock, and bond markets, *Journal of Econometrics*, 160, pp. 48-57.

Chkili, W., Hammoudeh, S. and Nguyen, D.K. (2014). Volatility forecasting and risk management for commodity markets in the presence of asymmetry and long memory, *Energy Economics*, 41, pp. 1-18.

Chou, R.Y. and Liu, N. (2010). The economic value of volatility timing using a range-based volatility model, *Journal of Economic Dynamics and Control*, 34, pp. 2288-2301.

Christiansen, C., Schmeling, M. and Schrimpf, A. (2012). A comprehensive look at financial volatility prediction by economic variables, *Journal of Applied Econometrics*, 27(6), 956-977.

Corsi, F. (2009). A simple approximate long-memory model of realized volatility, *Journal of Financial Econometrics*, 7(2), pp. 174-196.

Degiannakis, S. and Filis, G. (2017). Forecasting oil price realized volatility using information channels from other asset classes, *Journal of International Money and Finance*, 76, pp. 28-49.

Degiannakis, S. and Filis, G. (2022). Oil price volatility forecasts: What do investors need to know?, *Journal of International Money and Finance*, 123, 102594.

Degiannakis, S., Filis, G. and Hassani, H. (2018). Forecasting global stock market implied volatility indices, *Journal of Empirical Finance*, 46, pp. 111-129.

Delis, P., Degiannakis, S. and Filis, G. (2022). What matters when developing oil price volatility forecasting frameworks?, *Journal of Forecasting*, 41(2), 361-382.

Delis, P., Degiannakis, S. and Giannopoulos, K. (2023). What Should be Taken into Consideration when Forecasting Oil Implied Volatility Index?, *Energy Journal*, 44(5), pp. 231-250.

Elder, A. and Gannon, G. (1998). Evaluation of volatility forecasts in an economic value framework, *International Review of Financial Analysis*, 7(3), 221-236.

Elliott, G. and Timmermann, A. (2008). Economic Forecasting, *Journal of Economic Literature*, 46 (1), pp. 3-56.

Engle, R.F., Hong, C.H., Kane, A. and Noh, J. (1993). Arbitrage valuation of variance forecasts with simulated options, *Advances in Futures and Options Research*, 6, pp. 393–415.

Engle, R.F., Kane, A., Noh, J. (1996). Index-option pricing with stochastic volatility and the value of accurate variance forecasts, *Review of Derivatives Research*, 1 (2), pp. 139–157.

Fernandes, M., Medeiros, M.C. and Scharth, M. (2014). Modeling and predicting the CBOE market volatility index, *Journal of Banking and Finance*, 40, pp. 1-10.

Gong, X., Zhang, W., Zhao, Y. and Ye, X. (2023). Forecasting stock volatility with a large set of predictors: A new forecast combination method, *Journal of Forecasting*, *42*(7), 1622-1647.

Gu, S., Kelly, B. and Xiu, D. (2020). Empirical asset pricing via machine learning, *The Review of Financial Studies*, *33*(5), 2223-2273.

Hansen, L.P. and Hodrick, R.J. (1980). Forward exchange rates as optimal predictors of future spot rates: An econometric analysis, *Journal of Political Economy*, 88, pp. 829-853.

Hansen, P.R., Lunde, A. and Nason, J.M. (2011). The model confidence set, *Econometrica*, 79(2), 453-497.

Jiang, I-M., Hung, J-C. and Wang, C-S. (2014). Volatility forecasts: Do volatility estimators and evaluation methods matter?, *The Journal of Futures Markets*, 34(11), pp. 1077-1094.

Kambouroudis, D. S., McMillan, D.G. and Tsakou, K. (2021). Forecasting realized volatility: The role of implied volatility, leverage effect, overnight returns, and volatility of realized volatility, *Journal of Futures Markets*, 41(10), pp. 1618-1639.

Kim, C.J., Morley, J.C. and Nelson, C.R. (2004). Is there a positive relationship between stock market volatility and the equity premium?, *Journal of Money, Credit and Banking*, 339-360.

Kourtis, A., Markellos, R.N. and Symeonidis, L. (2016). An international comparison of implied, realized, and GARCH volatility forecasts, *The Journal of Futures Markets*, 36(12), pp. 1164-1193.

Liang, C., Tang, L., Li, Y. and Wei, Y. (2020a). Which sentiment index is more informative to forecast stock market volatility? Evidence from China, *International Review of Financial Analysis*, 71, 101552.

Liang, C., Wei, Y. and Zhang, Y. (2020b). Is implied volatility more informative for forecasting realized volatility: An international perspective, *Journal of Forecasting*, 39, pp. 1253-1276.

Ma, F., Wahab, M.I.M., Chevallier, J. and Li, Z. (2023). A tug of war of forecasting the US stock market volatility: Oil futures overnight versus intraday information, *Journal of Forecasting*, 42(1), 60-75.

Magner, N., Lavin, J.F., Valle, M. and Hardy, N. (2021). The predictive power of stock market's expectations volatility: A financial synchronization phenomenon, *PLoS ONE*, 16(5), pp. 1-21.

Noh, J., Engle, R.F. and Kane, A. (1994). Forecasting volatility and option prices of the S&P 500 index, *The Journal of Derivatives*, 2(1), pp. 17–30.

Patton, J.A. and Sheppard, K. (2015). Good volatility, bad volatility: Signed jumps and the persistence of volatility, *The Review of Economics and Statistics*, 97(3), pp. 683-697.

Qiao, G., Jiang, G. and Yang, J. (2022). VIX term structure forecasting: New evidence based on the realized semi-variances, *International Review of Financial Analysis*, 102199.

Richardson M. and Stock, J.H. (1989). Drawing inferences from statistics based on multilayer asset returns, *Journal of Financial Economics*, 25, pp. 323-348.

Salisu, A.A., Demirer, R. and Gupta, R. (2023). Policy uncertainty and stock market volatility revisited: the predictive role of signal quality, *Journal of Forecasting*, 42(8), 2307-2321.

Son, B., Lee, Y., Park, S. and Lee, J. (2023). Forecasting global stock market volatility: The impact of volatility spillover index in spatial-temporal graph-based model, *Journal of Forecasting*, 42(7), 1539-1559.

Taylor, N. (2014). The economic value of volatility forecasts: A conditional approach, *Journal of Financial Econometrics*, 12 (3), pp. 433-478.

Welch, I. and Goyal, A. (2008). A comprehensive look at the empirical performance of equity premium prediction, *The Review of Financial Studies*, *21*(4), 1455-1508.

Ye, W., Xia, W., Wu, B. and Chen, P. (2022). Using implied volatility jumps for realized volatility forecasting: Evidence from the Chinese market, *International Review of Financial Analysis*, 83, 102277.

Zhang, C., Zhang, Y., Cucuringu, M. and Qian, Z. (2024). Volatility forecasting with machine learning and intraday commonality, *Journal of Financial Econometrics*, 22(2), 492-530.

TABLES

Table 1: Descriptive statistics of the implied volatility indices. Period: 10th of May, 2007 – 3rd of April, 2020.

	VIX	BPVIX	EUVIX	JYVIX	OVX	TYVIX	VDAX	VNIK	VSTOXX
Mean	19.471	9.976	10.072	10.911	36.486	6.040	20.955	24.799	23.121
Median	16.490	9.080	9.350	10.540	32.960	5.470	18.870	22.390	20.933
Maximum	82.690	30.860	29.670	36.330	190.080	16.390	86.010	91.450	87.513
Minimum	9.140	4.330	3.990	4.290	14.500	3.160	10.880	12.190	10.678
Std. Dev.	9.872	3.961	3.895	3.683	16.106	2.063	8.322	10.232	9.772
Skewness	2.627	2.192	1.425	1.484	3.305	1.533	2.504	2.518	2.144
Kurtosis	11.709	9.214	5.980	7.318	22.840	5.732	12.239	11.904	9.666
Jarque-Bera	13522.070	7560.272	2221.958	3587.945	57161.020	2204.357	14436.450	13678.310	8212.214
Probability (JB)	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000
Coefficient of Variation (CV)	0.507	0.397	0.387	0.338	0.441	0.342	0.397	0.413	0.423

Table 2: Percentage of long positions for VIX futures trading. Out-of-sample period 27th April, 2011 - 3rd April, 2020.

Forecast horizon	HAR-IV	HAR-IV-X							
(s-days ahead)		BPVIX	EUVIX	JYVIX	OVX	TYVIX	VDAX	VNIK	VSTOXX
1	0.6231	0.6010	0.6075	0.6581	0.6567	0.6341	0.6539	0.6355	0.6659
5	0.6353	0.5976	0.6353	0.6679	0.6826	0.6367	0.6748	0.6527	0.6977
10	0.6492	0.6051	0.6400	0.6602	0.6961	0.6318	0.6970	0.6713	0.7121
22	0.6515	0.6129	0.6267	0.6713	0.6920	0.6299	0.7200	0.6786	0.7292
44	0.6751	0.6182	0.6315	0.6797	0.7058	0.6338	0.7242	0.6985	0.7306
66	0.6903	0.6257	0.6386	0.6922	0.7114	0.6473	0.7215	0.7036	0.7311

Note: The percentage of short positions is $1 - \%$ of long positions.

Table 3: Percentage of long positions for S&P500 futures trading. Out-of-sample period 27th April, 2011 - 3rd April, 2020.

Forecast horizon	HAR-IV	HAR-IV-X							
(s-days ahead)		BPVIX	EUVIX	JYVIX	OVX	TYVIX	VDAX	VNIK	VSTOXX
1	0.3769	0.3990	0.3925	0.3419	0.3433	0.3659	0.3461	0.3645	0.3341
5	0.3647	0.4024	0.3647	0.3321	0.3174	0.3633	0.3252	0.3473	0.3023
10	0.3508	0.3949	0.3600	0.3398	0.3039	0.3682	0.3030	0.3287	0.2879
22	0.3485	0.3871	0.3733	0.3287	0.3080	0.3701	0.2800	0.3214	0.2708
44	0.3249	0.3818	0.3685	0.3203	0.2942	0.3662	0.2758	0.3015	0.2694
66	0.3097	0.3743	0.3614	0.3078	0.2886	0.3527	0.2785	0.2964	0.2689

Note: The percentage of short positions is $1 - \%$ of long positions.

Table 4: Trading profits from VIX futures. Out-of-sample period 27th April, 2011 - 3rd April, 2020.

Forecast horizon	HAR-IV	HAR-IV-X							
(*s*-days ahead)		BPVIX	EUVIX	JYVIX	OVX	TYVIX	VDAX	VNIK	VSTOXX
1	790.00%	442.61%	661.61%	581.31%	663.37%	786.25%	540.07%	511.69%	511.63%
5	888.89%	584.58%	*861.44%*	704.78%	*896.17%*	866.28%	673.25%	694.98%	557.66%
10	867.44%	845.26%	827.09%	613.10%	825.50%	879.53%	850.95%	808.63%	624.22%
22	*917.41%*	*968.67%*	693.91%	684.79%	759.45%	***988.94%***	754.66%	741.86%	682.11%
44	837.90%	859.89%	673.03%	*836.86%*	719.10%	904.70%	834.80%	*824.78%*	772.90%
66	754.20%	782.90%	723.49%	714.66%	744.00%	935.24%	*856.62%*	759.69%	*883.58%*

Note: Numbers in *italics* denote highest profit levels per forecasting model. **Bold** numbers denote the models included in the confidence set of the models with the highest profits across all horizons.

Table 5: Trading profits from S&P500 futures. Out-of-sample period 27th April, 2011 - 3rd April, 2020.

Forecast horizon	HAR-IV	HAR-IV-X							
(*s*-days ahead)		BPVIX	EUVIX	JYVIX	OVX	TYVIX	VDAX	VNIK	VSTOXX
1	17.75%	10.24%	-11.65%	-4.11%	-18.97%	32.36%	-13.95%	-15.89%	-19.71%
5	35.12%	33.94%	*46.21%*	12.30%	*25.74%*	61.38%	-1.63%	3.62%	7.19%
10	*35.94%*	*39.05%*	33.25%	-3.23%	2.43%	63.48%	9.14%	*25.50%*	3.48%
22	25.78%	19.18%	16.65%	6.06%	-9.85%	***66.18%***	3.78%	17.92%	3.32%
44	9.12%	27.89%	25.57%	*12.83%*	-26.43%	45.99%	*14.37%*	12.89%	8.73%
66	-0.48%	12.75%	37.35%	3.62%	-21.70%	49.22%	11.17%	5.97%	*19.68%*

Note: Numbers in *italics* denote highest profit levels per forecasting model. **Bold** numbers denote the models included in the confidence set of the models with the highest profits across all horizons.

Table 6: Statistical Loss Function: RMSE. Out-of-sample period 27th April, 2011 - 3rd April, 2020.									
Forecast horizon	HAR-IV	HAR-IV-X							
(*s*-days ahead)		BPVIX	EUVIX	JYVIX	OVX	TYVIX	VDAX	VNIK	VSTOXX
1	**1.552***	1.556	1.554	1.555	1.554	1.554	1.555	1.554	1.558
5	**2.847***	2.859	2.860	2.861	2.869	2.848	2.872	2.855	2.882
10	3.400	3.415	3.426	3.411	3.445	**3.385***	3.434	3.413	3.454
22	4.115	4.147	4.150	4.098	4.205	**4.048***	4.157	4.149	4.182
44	4.833	4.738	4.826	4.727	4.977	**4.625***	4.895	4.914	4.869
66	7.020	6.834	6.894	**6.760***	7.069	6.770	7.049	7.118	7.028

Note: Numbers in **bold** denote lowest loss function per forecasting horizon.
* denotes that the model is included in the confidence set of the models with the lowest values in the RMSE loss function.

Table 7: Statistical Loss Function: MAE. Out-of-sample period 27th April, 2011 - 3rd April, 2020.									
Forecast horizon	HAR-IV	HAR-IV-X							
(*s*-days ahead)		BPVIX	EUVIX	JYVIX	OVX	TYVIX	VDAX	VNIK	VSTOXX
1	**0.955***	0.958	0.957	0.961	0.958	0.958	0.959	0.956	0.961
5	**1.861***	1.871	1.874	1.896	1.894	1.873	1.880	1.866	1.889
10	2.297*	**2.295***	2.318	2.340	2.366	2.296*	2.328	2.316	2.353
22	2.899	2.899	2.940	2.916	3.037	**2.841***	2.944	2.919	2.979
44	3.527	3.424	3.512	3.519	3.706	**3.368***	3.583	3.533	3.552
66	4.274	4.044	4.172	4.244	4.417	**4.018***	4.302	4.285	4.312

Note: Numbers in **bold** denote lowest loss function per forecasting horizon.
* denotes that the model is included in the confidence set of the models with the lowest values in the MAE loss function.

FIGURES

Figure 1: Evolution of the implied volatility indices. Period: 10th May, 2007 – 3rd April, 2020

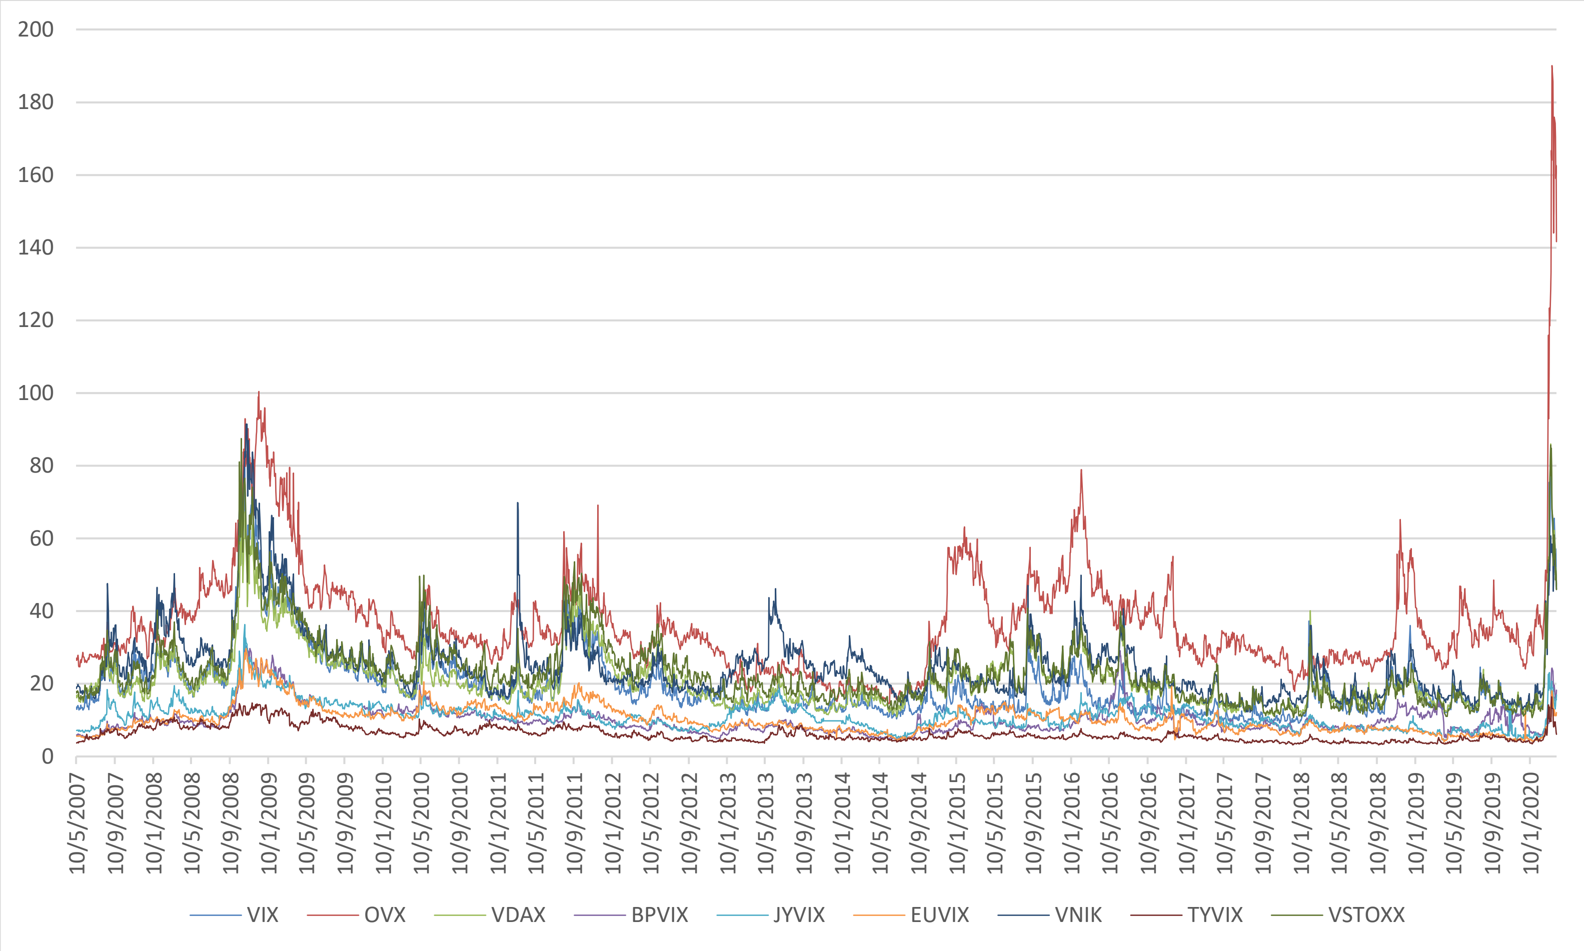

Figure 2: Total trading profits from VIX futures. Out-of-sample period 27th April, 2011 - 3rd April, 2020.

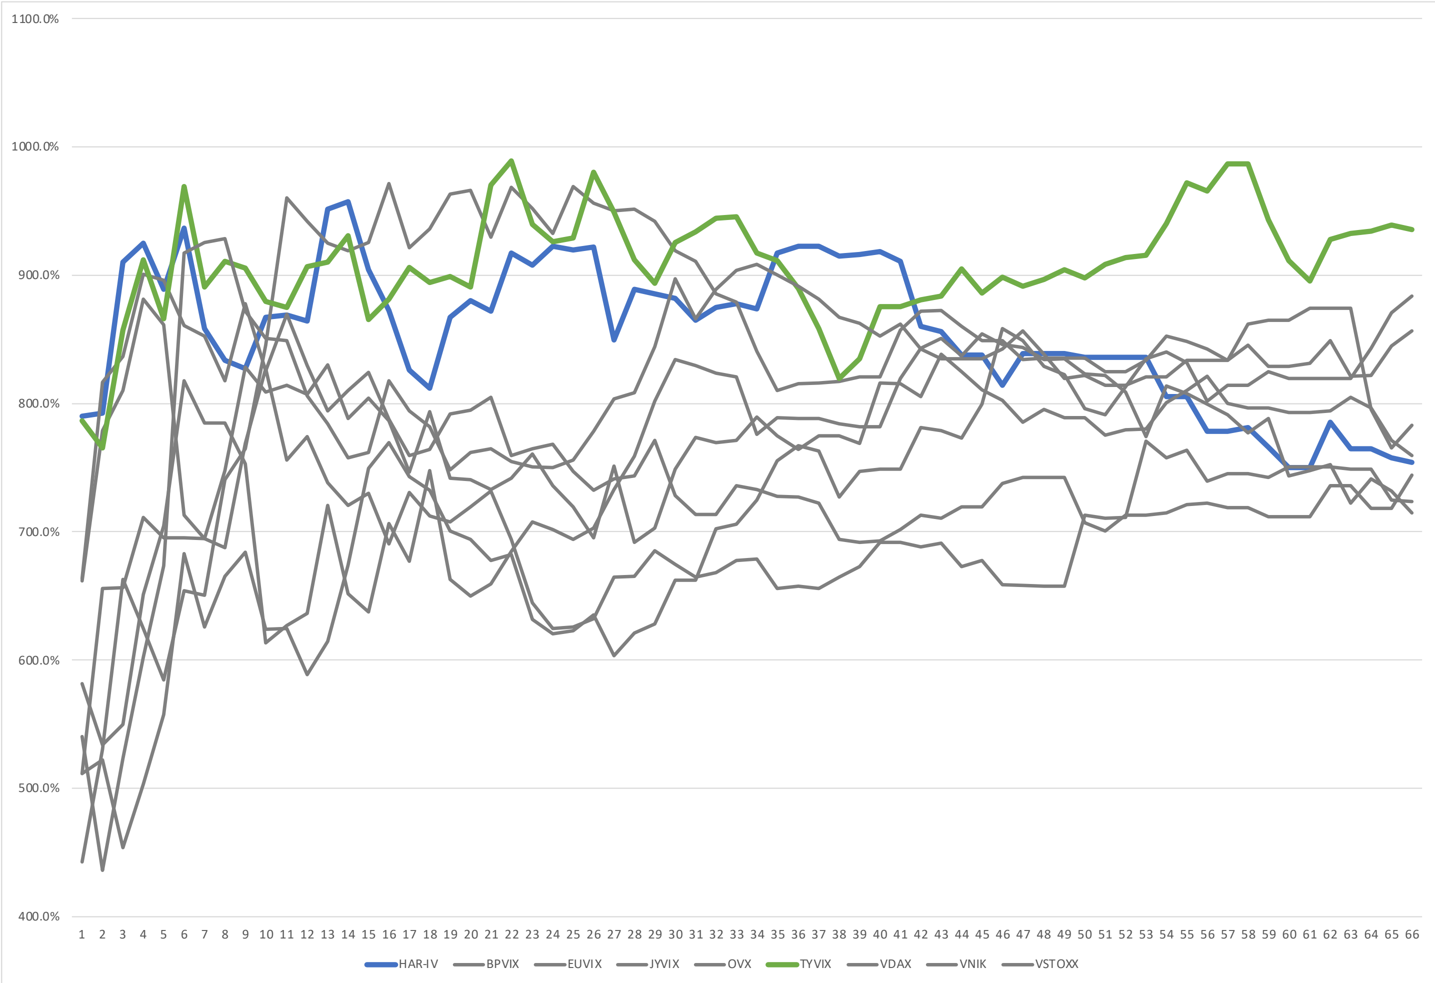

Note: We only highlight in colour the cumulative profits of the HAR-IV and HAR-IV-TYVIX to allow easier comparison.

Figure 3: Total trading profits from S&P500 futures. Out-of-sample period 27th April, 2011 - 3rd April, 2020.

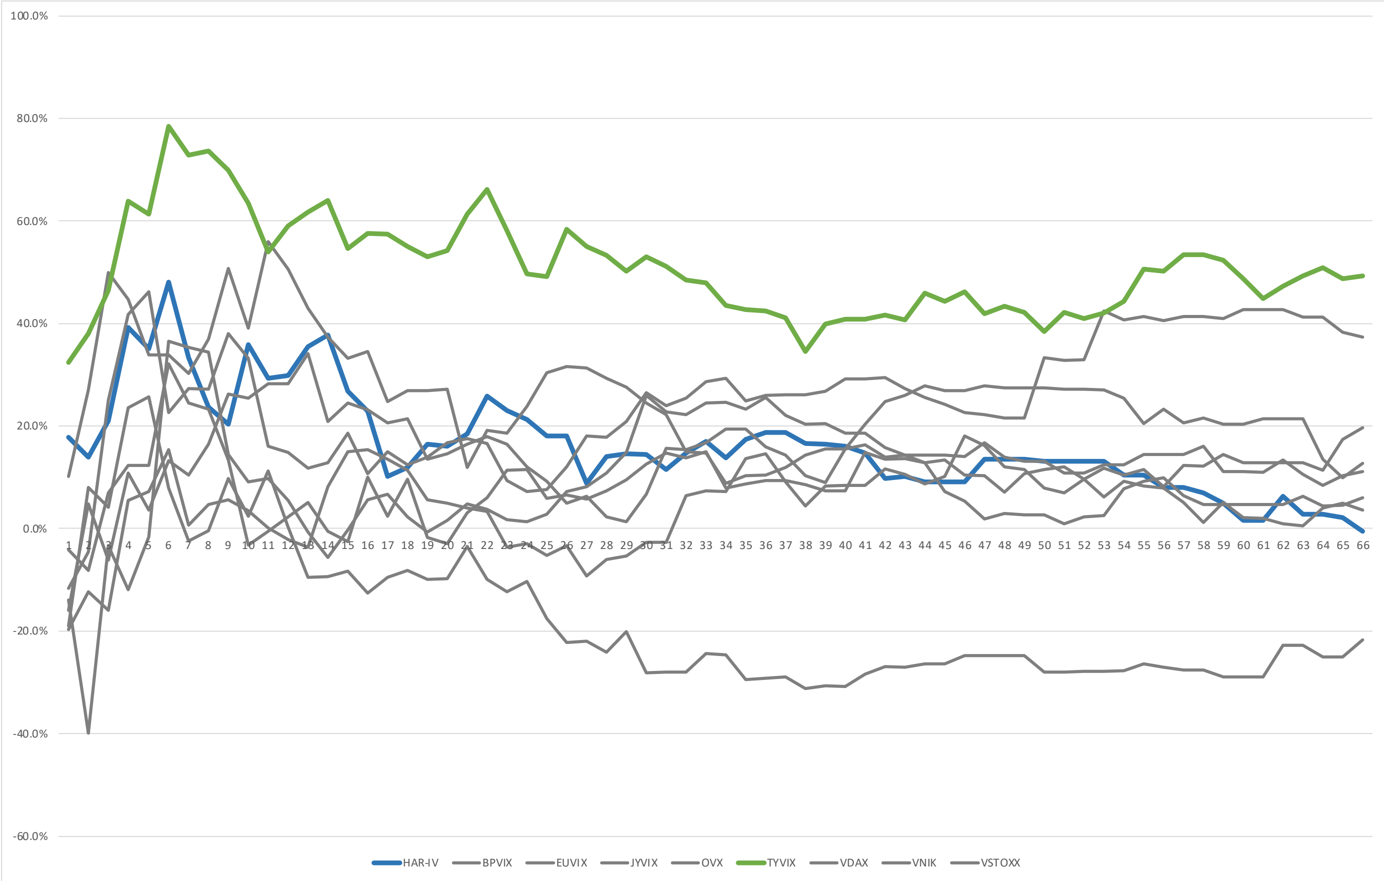

Note: We only highlight in colour the cumulative profits of the HAR-IV and HAR-IV-TYVIX to allow easier comparison.

Figure 4: Trading profits from VIX and S&P500 futures over the out-of-sample period: 27th April, 2011 - 3rd April, 2020.

VIX futures

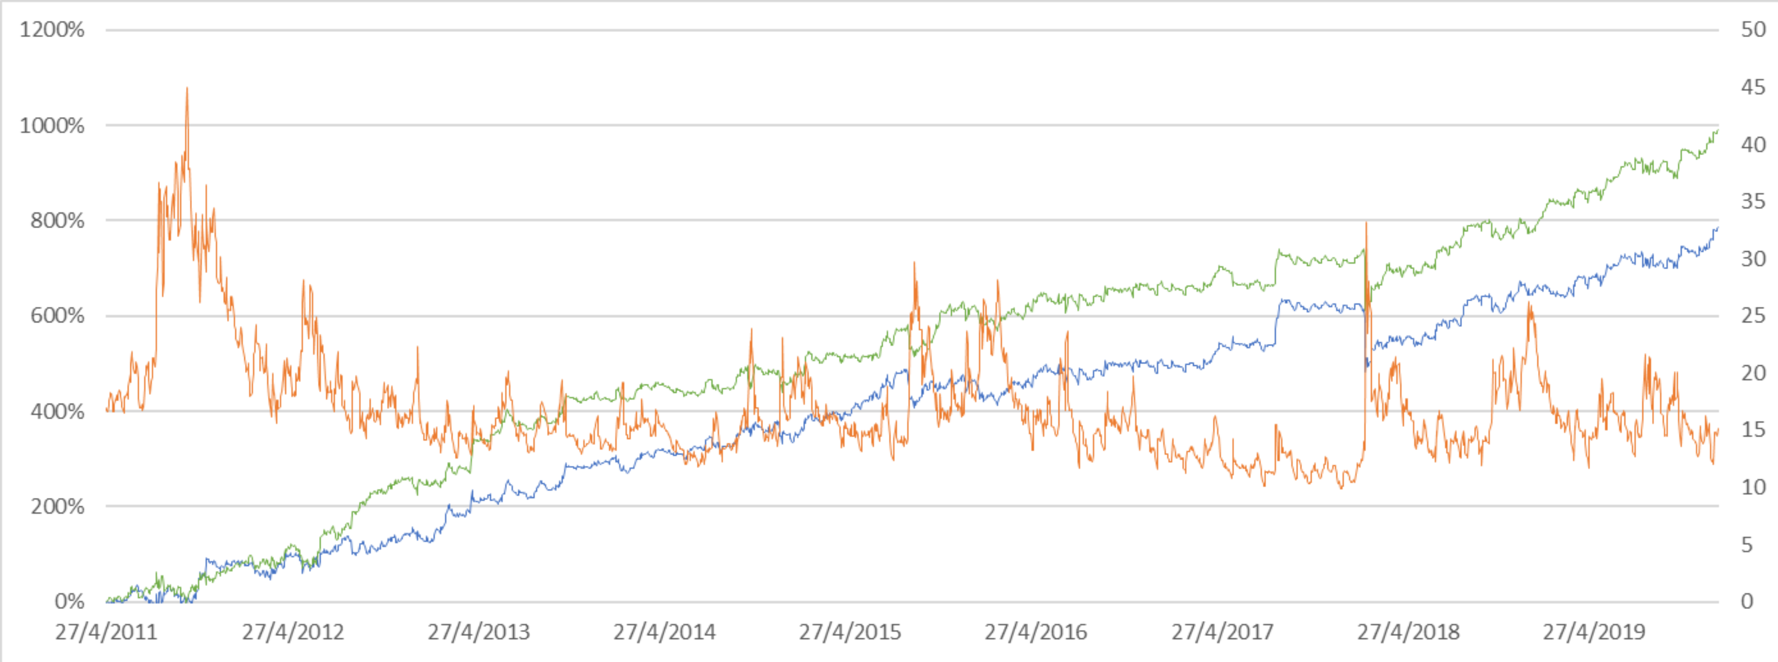

S&P500 futures

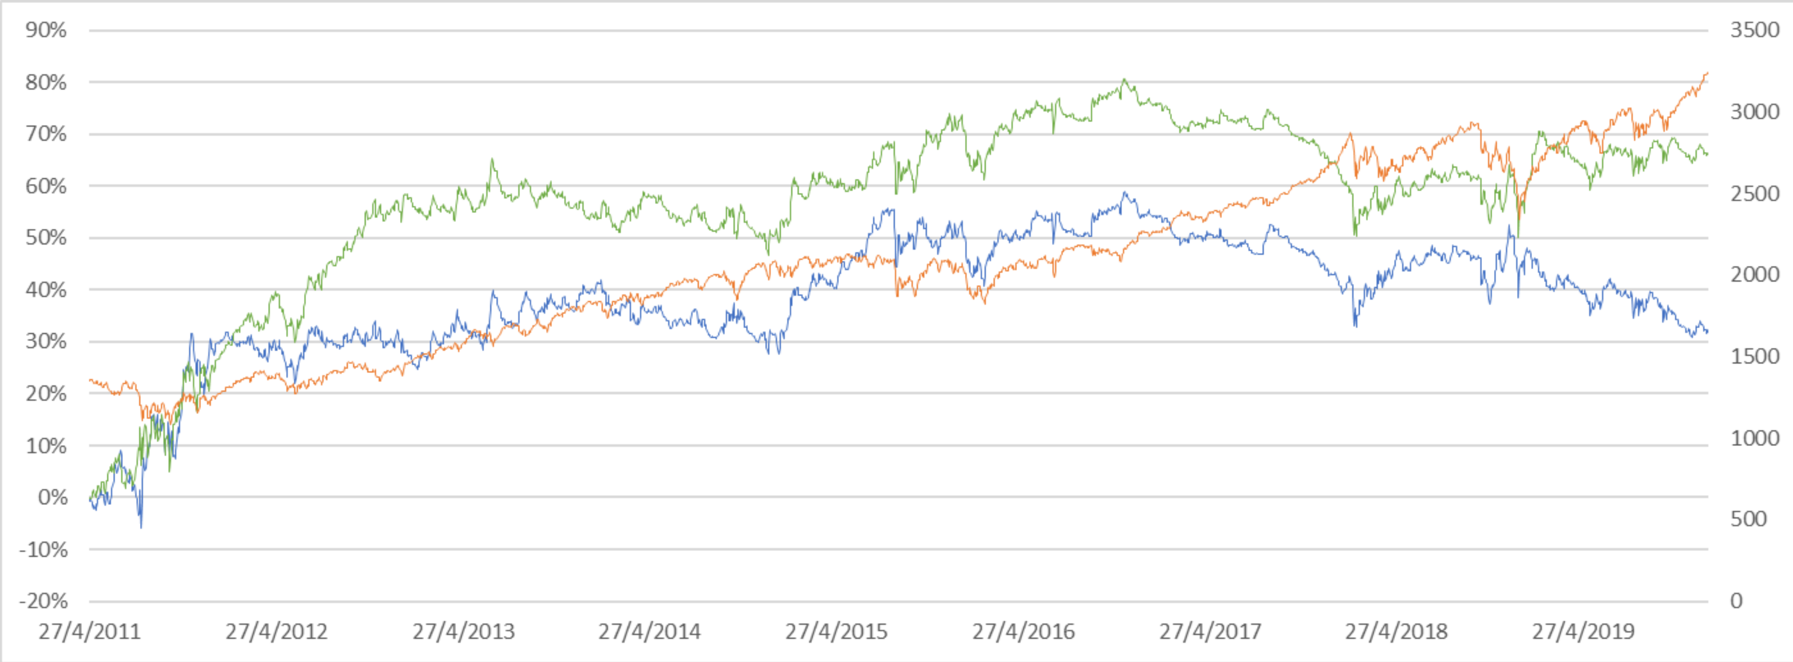

Note: The LHS axis shows the level of trading profits, whereas the RHS axis show the VIX (top panel) and S&P500 (bottom panel) futures values. The orange line depicts the VIX (top panel) and S&P500 (bottom panel) futures values over the out-of-sample period, whereas, the blue and green lines show the 1-day trading profits from the 1-day ahead and 22-days ahead forecasts from the HAR-IV-TYVIX, respectively.

Figure 5: Trading strategy performance. Sharpe, Sortino and Calmar ratios for S&P500 and VIX futures.

Sharpe ratio for S&P500 (left) and VIX (right) futures

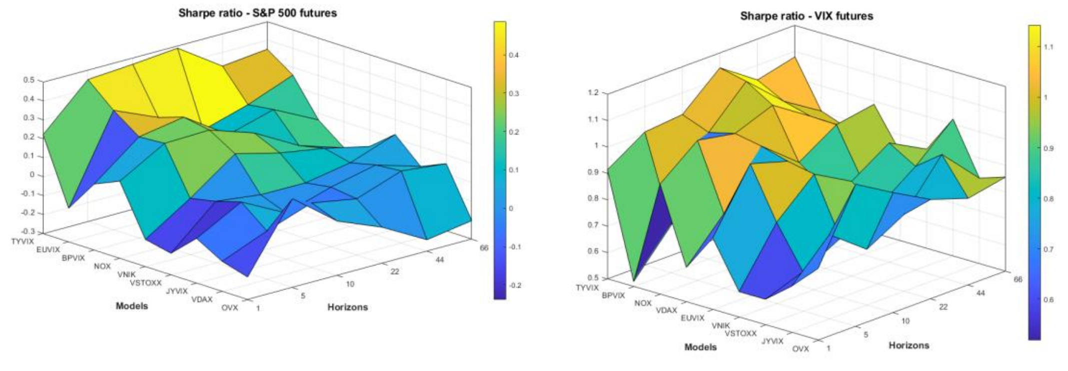

Sortino ratio for S&P500 (left) and VIX (right) futures

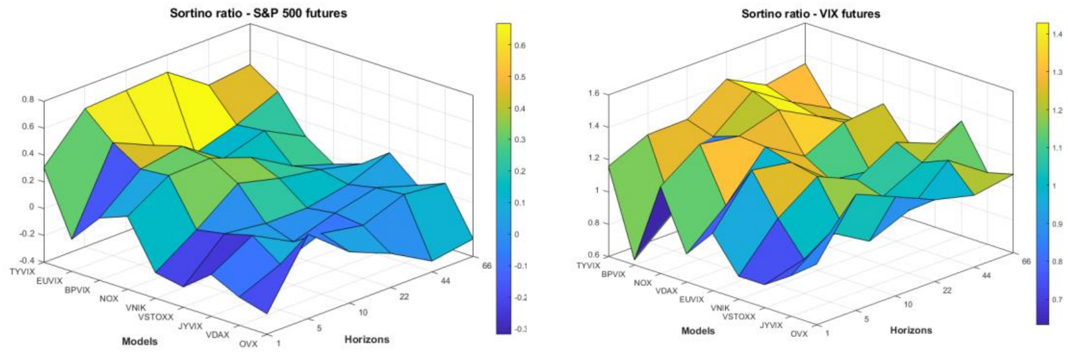

Calmar ratio for S&P500 (left) and VIX (right) futures

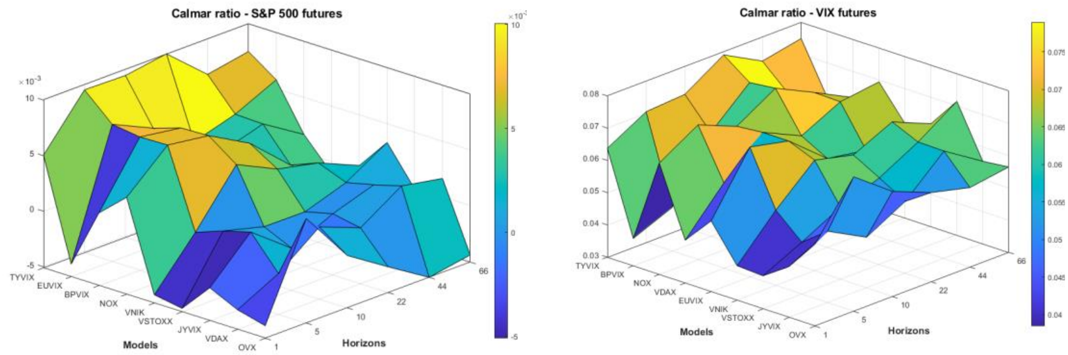

Figure 6: Trading profits from VIX and S&P500 futures over the post out-of-sample period: 6th April, 2020 - 18th February, 2022.

VIX futures

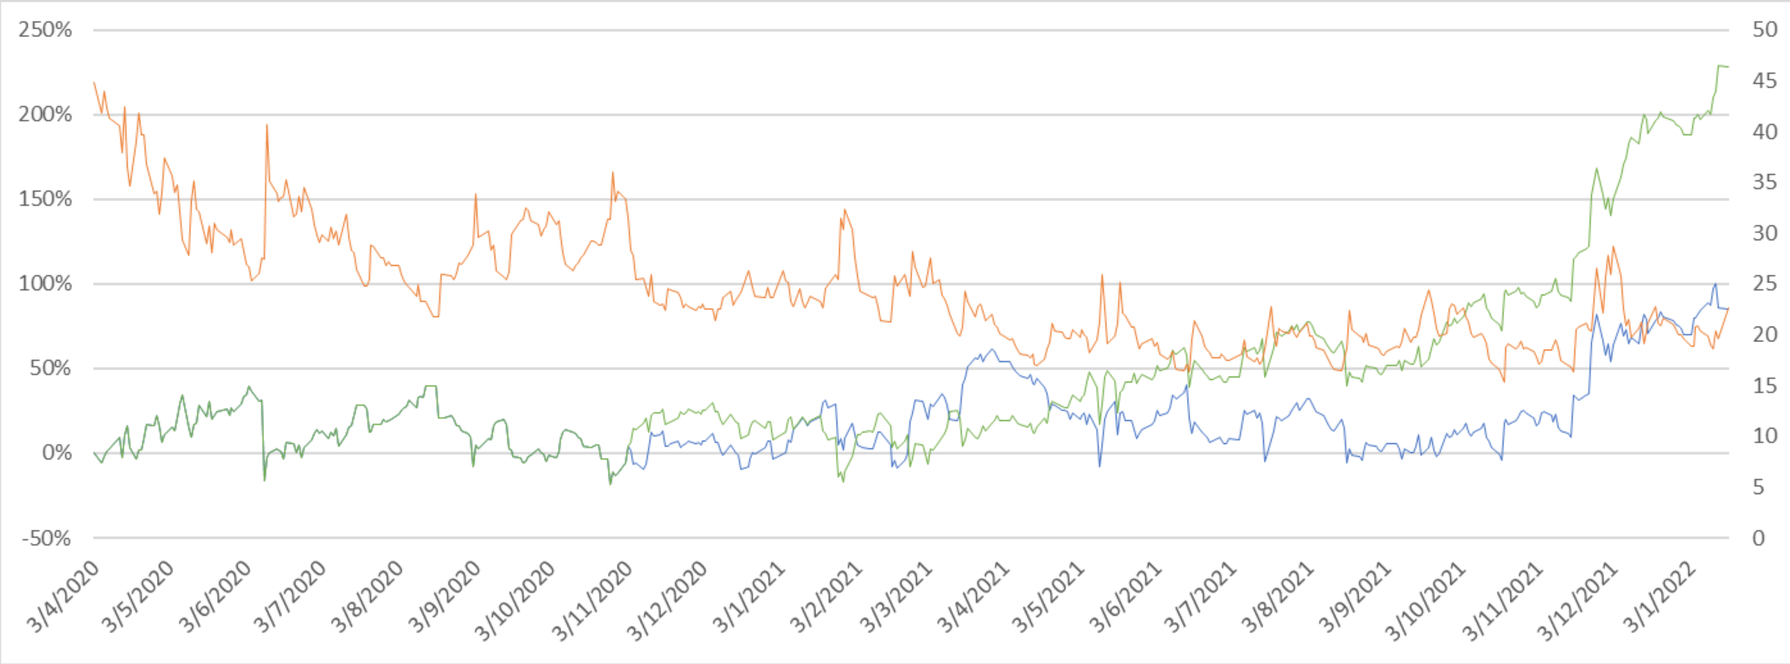

S&P500 futures

Note: The LHS axis shows the level of trading profits, whereas the RHS axis show the VIX (top panel) and S&P500 (bottom panel) futures values. The orange line depicts the VIX (top panel) and S&P500 (bottom panel) futures values over the out-of-sample period, whereas, the blue and green lines show the 1-day trading profits from the 1-day ahead and 22-days ahead forecasts from the HAR-IV, respectively. The blue and green lines coincide for the first period of the post out-of-sample period.

APPENDIX

Table A1: Trading profits from VIX futures. Fixed window length of 750 observations.

Forecast horizon	HAR-IV	HAR-IV-X							
(*s*-days ahead)		BPVIX	EUVIX	JYVIX	OVX	TYVIX	VDAX	VNIK	VSTOXX
1	650.98%	842.91%	746.93%	328.61%	593.39%	637.70%	466.75%	448.90%	574.57%
5	674.43%	534.42%	787.00%	492.25%	715.09%	836.13%	503.79%	663.53%	445.27%
10	706.45%	640.55%	801.91%	592.85%	749.37%	813.82%	600.86%	551.04%	498.71%
22	707.02%	667.11%	758.26%	804.94%	614.20%	***919.70%***	727.02%	550.51%	599.20%
44	713.08%	884.88%	723.23%	677.87%	638.87%	910.89%	621.94%	758.70%	415.46%
66	677.28%	832.27%	603.53%	696.76%	581.71%	626.84%	737.69%	736.75%	489.89%

Note: Numbers in *italics* denote highest profit levels per forecasting model. **Bold** numbers denote the model with the highest profits across all horizons and models.

Table A2: Trading profits from S&P500 futures. Fixed window length of 750 observations.

Forecast horizon	HAR-IV	HAR-IV-X							
(*s*-days ahead)		BPVIX	EUVIX	JYVIX	OVX	TYVIX	VDAX	VNIK	VSTOXX
1	15.43%	39.22%	-8.90%	-28.25%	-5.74%	8.94%	-16.06%	-25.49%	1.75%
5	26.73%	16.96%	14.27%	13.80%	31.69%	63.94%	-6.20%	34.26%	-12.52%
10	27.58%	6.71%	32.22%	-11.50%	31.38%	58.35%	-18.94%	5.78%	-12.66%
22	20.17%	-1.89%	32.85%	27.61%	11.42%	***65.72%***	2.69%	7.70%	4.26%
44	15.21%	38.59%	38.00%	-0.61%	8.31%	52.18%	-5.17%	13.39%	-11.65%
66	9.40%	33.15%	26.11%	5.12%	13.63%	6.80%	10.45%	6.77%	-5.08%

Note: Numbers in *italics* denote highest profit levels per forecasting model. **Bold** numbers denote the model with the highest profits across all horizons and models.

Table A3: Trading profits from VIX futures. Fixed window length of 1250 observations.

Forecast horizon (*s*-days ahead)	HAR-IV	HAR-IV-X							
		BPVIX	EUVIX	JYVIX	OVX	TYVIX	VDAX	VNIK	VSTOXX
1	644.92%	704.61%	688.46%	683.09%	556.56%	594.35%	661.76%	735.41%	630.13%
5	722.22%	660.86%	763.42%	782.40%	764.83%	670.33%	797.25%	784.85%	705.76%
10	752.62%	702.34%	784.50%	732.49%	756.08%	876.30%	725.09%	763.92%	690.18%
22	768.15%	718.96%	692.44%	822.49%	931.76%	***948.12%***	868.77%	792.87%	677.52%
44	791.89%	785.03%	724.80%	906.81%	798.61%	963.75%	668.24%	744.60%	747.16%
66	685.12%	856.93%	669.61%	757.31%	761.23%	855.41%	708.17%	765.71%	782.17%

Note: Numbers in *italics* denote highest profit levels per forecasting model. **Bold** numbers denote the model with the highest profits across all horizons and models.

Table A4: Trading profits from S&P500 futures. Fixed window length of 1250 observations.

Forecast horizon (*s*-days ahead)	HAR-IV	HAR-IV-X							
		BPVIX	EUVIX	JYVIX	OVX	TYVIX	VDAX	VNIK	VSTOXX
1	-13.18%	22.90%	-19.46%	-5.90%	-29.98%	-6.52%	-31.26%	-13.40%	-30.98%
5	-14.03%	11.35%	1.21%	-2.08%	-2.23%	-3.18%	-26.15%	-4.18%	-19.75%
10	-0.12%	10.18%	13.31%	-11.79%	-2.38%	34.02%	-17.49%	-4.71%	-21.62%
22	-12.42%	-8.75%	3.88%	-9.51%	-0.86%	***41.98%***	4.75%	5.77%	-23.02%
44	-13.95%	-1.37%	7.74%	-3.01%	-14.07%	27.51%	-38.23%	-8.21%	-20.43%
66	-27.95%	8.76%	-5.62%	-23.35%	-14.41%	5.97%	-33.03%	-6.63%	-14.87%

Note: Numbers in *italics* denote highest profit levels per forecasting model. **Bold** numbers denote the model with the highest profits across all horizons and models.

BANK OF GREECE WORKING PAPERS

317. Degiannakis, S. “The D-model for GDP nowcasting”, April 2023.
318. Degiannakis, S., G. Filis, G. Siourounis, L. Trapani, “Superkurtosis”, April 2023.
319. Dixon, H. T. Kosma, and P. Petroulas, “Explaining the endurance of price level differences in the euro area”, May 2023.
320. Kollintzas, T. and V. Vassilatos, “Implications of market and political power interactions for growth and the business cycle II: politico-economic equilibrium”, May 2023.
321. Bragoudakis, Z. and I. Krompas “Greek GDP forecasting using Bayesian multivariate models”, June 2023.
322. Degiannakis, S. and E. Kafousaki “Forecasting VIX: The illusion of forecast evaluation criteria”, June 2023.
323. Andreou C. P., S. Anyfantaki, C. Cabolis and K. Dellis, “Exploring country characteristics that encourage emissions reduction”, July 2023.
324. Dimakopoulou, V., Economides, G., Philippopoulos, A., and V. Vassilatos, “Can central banks do the unpleasant job that governments should do?”, December 2023.
325. Chrysanthakopoulos, C. and A. Tagkalakis, “The medium-term effects of fiscal policy rules”, January 2024.
326. Manou, K. and E. Papapetrou, “Does uncertainty matter for household consumption? A mean and a two tails approach”, February 2024.
327. Kakridis, A., “War, mobilization, and fiscal capacity: testing the bellicist theory in Greece, 1833-1939”, March 2024.
328. Mavrogiannis, C. and A. Tagkalakis, “From policy to capital: assessing the impact of structural reforms on gross capital inflows”, April 2024
329. Delis, P., S. Degiannakis, G. Filis, T. Palaskas and C. Stoforos, “Determinants of regional business cycle synchronization in Greece”, May 2024.
330. Sideris, D. and G. Pavlou, “Market power and profit margins in the Euro area countries in the post-pandemic period”, June 2024.
331. Kasimati, E. and N. Veraros, “The dry-bulk shipping market: a small econometric model”, September 2024.
332. Mermelas, G. and A. Tagkalakis, “Monetary policy transmission: the role of banking sector characteristics in the euro area”, November 2024.
333. Anastasiou, D., Pasiouras, F., Rizos, A., and A. Stratopoulou, “Do macroprudential policies make SMEs more-or-less discouraged to apply for a bank loan?”, December 2024.
334. Malliaropulos, D., Passari, E., and F. Petroulakis, “Unpacking commodity price fluctuations: reading the news to understand inflation”, December 2024
335. Degiannakis, S. and E. Kafousaki, “Disaggregating VIX”, January 2025