International Journal of Forecasting 32 (2016) 695–715

Contents lists available at ScienceDirect

International Journal of Forecasting

journal homepage: www.elsevier.com/locate/ijforecast

Combining nearest neighbor predictions and model-based predictions of realized variance: Does it pay?

Julián Andrada-Félix [a], Fernando Fernández-Rodríguez [a], Ana-Maria Fuertes [b,*]

[a] Universidad de Las Palmas de Gran Canaria, Department of Quantitative Methods in Economics and Management, Campus de Tafira 35017, Spain

[b] Cass Business School, City University London, 106 Bunhill Row, London ECIY 8TZ, UK

ARTICLE INFO

Keywords:
Realized volatility
Volatility forecasting
Non-parametric forecasts
Nearest neighbor
Long-memory models
Forecast combination
Straddles
Options trading

ABSTRACT

The increasing availability of intraday financial data has led to improvements in daily volatility forecasting through the use of long-memory models of realized volatility. This paper demonstrates the merit of the non-parametric nearest neighbor (NN) approach for S&P 100 realized variance forecasting. The NN approach is appealing a priori because, unlike model-based methods, it can reproduce complex dynamic dependencies, while largely avoiding misspecification and parameter estimation uncertainty. We evaluate the forecasts through straddle trading profitability metrics and using conventional statistical accuracy criteria. The ranking of individual forecasts confirms that there is not a one-to-one mapping between statistical accuracy and profitability. In turbulent markets, the NN forecasts lead to higher risk-adjusted profitability levels, even though the model-based forecasts are superior statistically. A directional combination of NN and model-based forecasts is more profitable than any of the individual forecasts, in both calm and turbulent market conditions.

1. Introduction

One decision that confronts both academic researchers and practitioners who face the task of financial time series forecasting is whether to use a non-parametric or a model-based method. Non-parametric methods are attractive for variables with complex dynamics that otherwise would require heavily parameterized models. A combination of forecasts from non-parametric methods and time series models provides protection against model misspecification and parameter estimation uncertainty, given the distinct way in which each exploits the information set (Timmermann, 2006). The goal of this paper is to demonstrate the merit of the non-parametric nearest neighbor (NN) method in the novel context of *realized volatility* prediction. The motivation for choosing the realized volatility as the target variable is threefold.

First, in contrast to asset returns, which are often portrayed as martingale difference series, the volatility displays persistence, and therefore should be predictable. Second, volatility forecasts are key inputs in financial applications such as derivatives pricing, risk management and portfolio allocation. Third, a paradigm shift has occurred in the volatility forecasting literature since the seminal papers by Andersen and Bollerslev (1998) and French, Schwert, and Stambaugh (1987), such that, instead of adopting daily GARCH or stochastic volatility models that treat the volatility as latent, many studies construct forecasts from long memory models fitted to daily realized

* Corresponding author.
E-mail addresses: julian.andrada@ulpgc.es (J. Andrada-Félix), fernando.fernandez@ulpgc.es (F. Fernández-Rodríguez), a.fuertes@city.ac.uk (A.-M. Fuertes).

http://dx.doi.org/10.1016/j.ijforecast.2015.10.004
0169-2070/

volatilities.[1] Recent studies have suggested a range of extensions of standard long memory models of the realized volatility for capturing various nonlinearities. However, little attention has been paid to the use of non-parametric forecasting methods for the realized volatility.

The non-parametric NN method that we consider in this paper is a "machine learning" tool that has been applied successfully in various pattern recognition studies in the areas of engineering and physics since its inception several decades ago, but is far less well known in finance, except for foreign exchange prediction (Arroyo & Mate, 2009; Fernández-Rodríguez, Sosvilla-Rivero, & Andrada-Félix, 1999). The merits of NN prediction in the context of the realized market volatility have not yet been studied. Our paper seeks to fill this gap. The most alluring aspect of the NN approach in this context is that it does not require the specification of a functional form for describing the dynamics of the realized volatility, which is potentially complex. One stylized fact of the realized volatility is long range dependence; however, it may also exhibit other features, such as news asymmetry, regime-switching dependence (with the 'switch' dictated by either observable or latent variables), and other nonlinear dependencies induced by market microstructure effects or measurement errors.

In essence, the NN approach identifies the local histories within the available time series of realized variances that are the most similar to the last observed history, then pools the subsequent observations of each history in order to build a prediction. The prediction can therefore be interpreted as a 'projection' from sequential local non-parametric regressions. Thus, NN forecasts can capture complex nonlinear dependences parsimoniously.

This paper contributes to the realized volatility forecasting literature by investigating whether the NN scheme can compete with long-memory models, and whether it is worthwhile to combine the forecasts from the two approaches. The evaluation tools include economic value measures and conventional statistical criteria. Our choice of an economic scenario is a *straddle* trading strategy that is informed by stock market realized variance predictions. The performance is assessed using conventional risk-adjusted profitability metrics in the context of Markowitz mean–variance analysis (e.g., Sharpe ratios), as well as assorted non-normality robust metrics. The evaluation accommodates a range of trading scenarios through the use of different noise-to-signal filtering rules, transaction cost levels, forecast horizons, and 'calm' versus 'turbulent' market conditions. As a byproduct, this comprehensive evaluation framework allows us to examine the degree of correspondence between the volatility forecast ranking implied by straddle-trading profitability measures and that arising from purely statistical criteria.

The study is based on five-minute intraday data on the S&P 100 stock index from January 1997 until November 2012.[2] The forecast evaluation is carried out over two distinct holdout (or trading) periods of similar lengths: the 2003–2007 period, which can be described as relatively 'calm' market conditions, and the 2008–2012 period, which largely reflects 'turmoil'. The start of the second period is an important landmark of the late 2000s global financial crisis, when various events shook the financial markets (e.g., Lehman Brothers bankruptcy).

We focus on the most popular of the time series models used previously for capturing long memory (and nonlinearities) in the realized volatility,. The combination approaches used are a unanimity rule or directional forecast combination, by which a trading signal is triggered only if the directions of the model-based and NN forecasts coincide; equally-weighted (EW) forecast combination; and OLS-weighted forecast combination.

The findings suggest that NN realized volatility forecasts are more effective for informing straddle trading than model-based forecasts during the post-2007 period. This confirms that the NN approach provides a 'shield' against the heightened misspecification and parameter estimation uncertainty experienced when the markets are in turmoil. Furthermore, the comparison of individual and combined forecasts reveals that the largest risk-adjusted profits in both calm and turbulent markets stem from the combination of NN and model-based predictions.

As a byproduct, the paper documents a mismatch between the statistical and economic rankings of realized volatility forecasts. The evidence adds to a body of literature which contends that statistically accurate predictions of the first or second moment of the distribution of asset returns do not necessarily imply profitability; see e.g. Bernales and Guidolin (2014), Cenesizoglu and Timmermann (2012), Gonçalves and Guidolin (2006), and Satchell and Timmermann (1995). This is because the loss function implicit in a (possibly simple) trading strategy is not necessarily captured well by the mean squared error or other statistical accuracy metrics.

The remainder of the paper is organized as follows. Section 2 briefly reviews the three strands of the literature that motivate our paper. Section 3 describes the dataset. Sections 4 and 5 present the methodologies for forecast construction and evaluation, respectively. Section 6 discusses the main empirical results. Finally, Section 7 concludes the paper.

2. Background literature

A burgeoning body of empirical literature that derives volatility forecasts from long-memory models of realized volatilities has emerged over the past two decades. This is motivated by the fact that long-range dependence

[1] The theory behind realized volatility estimators establishes that increasing the sampling frequency of asset returns makes it possible to obtain arbitrarily precise estimates of the daily volatility; that is, the daily volatility essentially becomes observable *ex post* (Andersen, Bollerslev, Diebold, & Labys, 2003; Barndorff-Nielsen & Shephard, 2002). Various studies have shown empirically that volatility forecast improvements can be obtained via long-memory models of the realized volatility; see e.g. Fuertes and Olmo (2012); Koopman, Jungbacker, and Hol (2005), and Pong, Shackleton, Taylor, and Xu (2004).

[2] Although S&P 100 options have been largely replaced by S&P 500 options in the empirical finance literature, the S&P 100 stock index is still used in practice for derivatives, e.g., OEX options.