Journal of Econometrics 192 (2016) 1–18

Contents lists available at ScienceDirect

Journal of Econometrics

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

Exploiting the errors: A simple approach for improved volatility forecasting

Tim Bollerslev[a,b,c], Andrew J. Patton[a,d,*], Rogier Quaedvlieg[e]

[a] *Department of Economics, Duke University, Durham, NC 27708, United States*
[b] *NBER, United States*
[c] *CREATES, Denmark*
[d] *Stern School of Business, New York University, United States*
[e] *Department of Finance, Maastricht University, Netherlands*

ARTICLE INFO

Article history:
Received 22 March 2015
Received in revised form
5 September 2015
Accepted 12 October 2015
Available online 10 November 2015

JEL classification:
C22
C51
C53
C58

Keywords:
Realized volatility
Forecasting
Measurement errors
HAR
HARQ

ABSTRACT

We propose a new family of easy-to-implement realized volatility based forecasting models. The models exploit the asymptotic theory for high-frequency realized volatility estimation to improve the accuracy of the forecasts. By allowing the parameters of the models to vary explicitly with the (estimated) degree of measurement error, the models exhibit stronger persistence, and in turn generate more responsive forecasts, when the measurement error is relatively low. Implementing the new class of models for the S&P 500 equity index and the individual constituents of the Dow Jones Industrial Average, we document significant improvements in the accuracy of the resulting forecasts compared to the forecasts from some of the most popular existing models that implicitly ignore the temporal variation in the magnitude of the realized volatility measurement errors.

1. Introduction

Volatility, and volatility forecasting in particular, plays a crucial role in asset pricing and risk management. Access to accurate volatility forecasts is of the utmost importance for many financial market practitioners and regulators. A long list of competing GARCH and stochastic volatility type formulations have been proposed in the literature for estimating and forecasting financial market volatility. The latent nature of volatility invariably complicates implementation of these models. The specific parametric models hitherto proposed in the literature generally also do not perform well when estimated directly with intraday data, which is now readily available for many financial assets. To help circumvent these complications and more effectively exploit the information inherent in high-frequency data, Andersen et al. (2003) suggested the use of reduced form time series forecasting models for the daily so-called realized volatilities constructed from the summation of the squared high-frequency intraday returns.[1]

Set against this background, we propose a new family of easy-to-implement volatility forecasting models. The models directly exploit the asymptotic theory for high-frequency realized volatility estimation by explicitly allowing the dynamic parameters of the models, and in turn the forecasts constructed from the models, to vary with the degree of estimation error in the realized volatility measures.

The realized volatility for most financial assets is a highly persistent process. Andersen et al. (2003) originally suggested the

* Correspondence to: Department of Economics, Duke University, 213 Social Sciences Building, Box 90097, Durham, NC 27708-0097, United States.
E-mail address: andrew.patton@duke.edu (A.J. Patton).

1 The use of realized volatility for accurately measuring the true latent integrated volatility was originally proposed by Andersen and Bollerslev (1998), and this approach has now become very popular for both measuring, modeling and forecasting volatility; see, e.g., the discussion and many references in the recent survey by Andersen et al. (2013).

http://dx.doi.org/10.1016/j.jeconom.2015.10.007
0304-4076/

use of fractionally integrated ARFIMA models for characterizing this strong dependency. However, the simple and easy-to-estimate approximate long-memory HAR (Heterogeneous AR) model of Corsi (2009) has arguably emerged as the preferred specification for realized volatility based forecasting. Empirically, the volatility forecasts constructed from the HAR model, and other related reduced-form time series models that treat the realized volatility as directly observable, generally perform much better than the forecasts from traditional parametric GARCH and stochastic volatility models.[2]

Under certain conditions, realized volatility (RV) is consistent (as the sampling frequency goes to zero) for the true latent volatility, however in any given finite sample it is, of course, subject to measurement error. As such, RV will be equal to the sum of two components: the true latent Integrated Volatility (IV) and a measurement error. The dynamic modeling of RV for the purposes of forecasting the true latent IV therefore suffers from a classical errors-in-variables problem. In most situations this leads to what is known as an attenuation bias, with the directly observable RV process being less persistent than the latent IV process. The degree to which this occurs obviously depends on the magnitude of the measurement errors; the greater the variance of the errors, the less persistent the observed process.[3]

Standard approaches for dealing with errors-in-variables problems treat the variance of the measurement error as constant through time.[4] In contrast, we explicitly take into account the temporal variation in the errors when modeling the realized volatility, building on the asymptotic distribution theory for the realized volatility measure developed by Barndorff-Nielsen and Shephard (2002). Intuitively, on days when the variance of the measurement error is small, the daily RV provides a stronger signal for next day's volatility than on days when the variance is large (with the opposite holding when the measurement error is large). Our new family of models exploits this heteroskedasticity in the error, by allowing for time-varying autoregressive parameters that are high when the variance of the realized volatility error is low, and adjusted downward on days when the variance is high and the signal is weak. Our adjustments are straightforward to implement and can easily be tailored to any autoregressive specification for RV. For concreteness, however, we focus our main discussion on the adaptation to the popular HAR model, which we dub the HARQ model. But, in our empirical investigation we also consider a number of other specifications and variations of the basic HARQ model.

Our empirical analysis relies on high-frequency data from 1997–2013 and corresponding realized volatility measures for the S&P 500 index and the individual constituents of Dow Jones Industrial Average. By explicitly incorporating the time-varying variance of the measurement errors into the parameterization of the model, the estimated HARQ models exhibit more persistence in "normal times" and quicker mean reversion in "erratic times" compared to the standard HAR model with constant autoregressive parameters.[5] Applying the HARQ model in an extensive out-of-sample forecast comparison, we document significant improvements in the accuracy of the forecasts compared to the forecasts from a challenging set of commonly used benchmark models. Interestingly, the forecasts from the HARQ models are not just improved in times when the right-hand side RVs are very noisy, and thus contain little relevant information, but also during tranquil times, when the forecasts benefit from the higher persistence afforded by the new models. Consistent with the basic intuition, the HARQ type models also offer the largest gains over the standard models for the assets for which the temporal variation in the magnitudes of the measurement errors are the highest.

The existing literature related to the dynamic modeling of RV and RV-based forecasting has largely ignored the issue of measurement errors, and when it has been considered, the errors have typically been treated as homoskedastic. Andersen et al. (2011), for instance, advocate the use of ARMA models as a simple way to account for measurement errors, while Asai et al. (2012) estimate a series of state-space models for the observable RV and the latent IV state variable with homoskedastic innovations. The approach for estimating stochastic volatility models based on realized volatility measures developed by Dobrev and Szerszen (2010) does incorporate the variance of the realized volatility error into the estimation of the models, but the parameters of the estimated models are assumed to be constant, and as such the dynamic dependencies and the forecasts from the models are not directly affected by the temporal variation in the size of the measurement errors. The motivation for the new family of HARQ models also bears some resemblance to the GMM estimation framework recently developed by Li and Xiu (2013). The idea of the paper is also related to the work of Bandi et al. (2013), who advocate the use of an "optimal", and possibly time-varying, sampling frequency when implementing RV measures, as a way to account for heteroskedasticity in the market microstructure "noise". In a similar vein, Shephard and Xiu (2014) interpret the magnitude of the parameter estimates associated with different RV measures in a GARCH-X model as indirect signals about the quality of the different measures: the lower the parameter estimate, the less smoothing, and the more accurate and informative the specific RV measure.

The rest of the paper is structured as follows. Section 2 provides the theoretical motivation for the new class of models, together with the results from a small scale simulation study designed to illustrate the workings of the models. Section 3 reports the results from an empirical application of the basic HARQ model for forecasting the volatility of the S&P 500 index and the individual constituents of the Dow Jones Industrial Average. Section 4 provides a series of robustness checks and extensions of the basic HARQ model. Section 5 concludes.

2. Realized volatility-based forecasting and measurement errors

2.1. Realized variance and high-frequency distribution theory

To convey the main idea, consider a single asset for which the price process P_t is determined by the stochastic differential

[2] Andersen et al. (2004) and Sizova (2011) show how minor model misspecification can adversely affect the forecasts from tightly parameterized volatility models, thus providing a theoretical explanation for this superior reduced-form forecast performance.

[3] Alternative realized volatility estimators have been developed by Barndorff-Nielsen et al. (2008), Zhang et al. (2005) and Jacod et al. (2009) among others. Forecasting in the presence of microstructure "noise" has also been studied by Aït-Sahalia and Mancini (2008), Andersen et al. (2011), Ghysels and Sinko (2011) and Bandi et al. (2013). The analysis below effectively abstracts from these complications, by considering a coarse five-minute sampling frequency and using simple RV. We consider some of these alternative estimators in Section 4.1.

[4] General results for the estimation of autoregressive processes with measurement error are discussed in Staudenmayer and Buonaccorsi (2005). Hansen and Lunde (2014) have also recently advocated the use of standard instrumental variable techniques for estimating the persistence of the latent IV process, with the resulting estimates being significantly more persistent than the estimates for the directly observable RV process.

[5] The persistence of the estimated HARQ models at average values for the measurement errors is very similar to the unconditional estimates based on Hansen and Lunde (2014), and as such also much higher than the persistence of the standard HAR models. We discuss this further below.