Journal of Econometrics 150 (2009) 119–138

Contents lists available at ScienceDirect

Journal of Econometrics

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

Predictive density estimators for daily volatility based on the use of realized measures[☆]

Valentina Corradi [a,*], Walter Distaso [b], Norman R. Swanson [c]

[a] *University of Warwick, Department of Economics, Coventry, CV4 7AL, UK*

[b] *Imperial College, Business School, South Kensington Campus, London SW7 2AZ, UK*

[c] *Department of Economics, Rutgers University, 75 Hamilton Street, New Brunswick, NJ 08901, USA*

ARTICLE INFO

Article history:
Available online 8 January 2009

JEL classification:
C22
C53
C14

Keywords:
Diffusions
Integrated volatility
Realized volatility measures
Kernels
Microstructure noise

ABSTRACT

The main objective of this paper is to propose a feasible, model free estimator of the predictive density of integrated volatility. In this sense, we extend recent papers by Andersen et al. [Andersen, T.G., Bollerslev, T., Diebold, F.X., Labys, P., 2003. Modelling and forecasting realized volatility. Econometrica 71, 579–626], and by Andersen et al. [Andersen, T.G., Bollerslev, T., Meddahi, N., 2004. Analytic evaluation of volatility forecasts. International Economic Review 45, 1079–1110; Andersen, T.G., Bollerslev, T., Meddahi, N., 2005. Correcting the errors: Volatility forecast evaluation using high frequency data and realized volatilities. Econometrica 73, 279–296], who address the issue of pointwise prediction of volatility via ARMA models, based on the use of realized volatility. Our approach is to use a realized volatility measure to construct a non-parametric (kernel) estimator of the predictive density of daily volatility. We show that, by choosing an appropriate realized measure, one can achieve consistent estimation, even in the presence of jumps and microstructure noise in prices. More precisely, we establish that four well known realized measures, i.e. realized volatility, bipower variation, and two measures robust to microstructure noise, satisfy the conditions required for the uniform consistency of our estimator. Furthermore, we outline an alternative simulation based approach to predictive density construction. Finally, we carry out a simulation experiment in order to assess the accuracy of our estimators, and provide an empirical illustration that underscores the importance of using microstructure robust measures when using high frequency data.

1. Introduction

In a recent paper, Andersen et al. (2003) suggest a novel, model free, approach for forecasting daily volatility. They advocate the use of simple, reduced form time series models for realized volatility, where the latter is constructed by summing up intradaily squared returns. The predictive ability of a given model is measured via the R^2 from the autoregressive or ARMA models constructed using (the log of) realized volatility. Their findings suggest that these ARMA based forecasts for realized volatility outperform most of the volatility models commonly used by practitioners, such as different varieties of GARCH models, for example. The rationale behind their approach is that, as the time interval between successive observations shrinks, realized volatility converges to the "true" daily volatility, whenever the underlying asset price is a continuous semimartingale. Although tick by tick and ultra high frequency data are now available, they are often contaminated by microstructure noise; therefore, in order to account for this potential problem, volatility has typically been constructed using 5 min interval returns, say, or even lower frequency observations.

Hence, these reduced form time series forecasts for realized volatility imply a loss in efficiency relative to the infeasible optimal forecasts for the daily volatility process, based on the entire volatility path. For the class of eigenfunction stochastic volatility models of Meddahi (2001), an analytical expression for such loss in efficiency is provided by Andersen et al. (2004). In particular, they show that the error associated with realized volatility induces a downward bias in the estimated degree of predictability obtained

☆ We thank the editors, Yongmiao Hong and Chung-Ming Kuan, two anonymous referees, Yacine Aït-Sahalia, Tim Bollerslev, Christian Gourieroux, Nour Meddahi, Victoria Zinde Walsh and participants at the Cirano-Cireq Conference on "Forecasting in Macroeconomics and Finance", the SETA conference on "Recent Developments in Financial Econometrics", the conference "Changing Structures in International and Financial Markets and the Effects of Financial Decision Making" in Venice, for helpful comments and suggestions. Corradi and Distaso gratefully acknowledge ESRC grant RES-000-23-0006, and Swanson acknowledges financial support from a Rutgers University Research Council grant.

* Corresponding author.
E-mail addresses: v.corradi@warwick.ac.uk (V. Corradi), w.distaso@imperial.ac.uk (W. Distaso), nswanson@econ.rutgers.edu (N.R. Swanson).

0304-4076/$ – see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.jeconom.2008.12.015

via the R^2 approach mentioned above. To overcome this issue, Andersen et al. (2005) develop a general, model free, feasible procedure to compute the adjusted R^2 used in model evaluation. Galbraith and Zinde Walsh (2006) use realized volatility measures to obtain consistent Least Square and Least Absolute Deviation deviation estimators of GARCH parameters, regardless of the implied volatility measurement error. More recently, Andersen et al. (2006), Aït-Sahalia and Mancini (in press) and Ghysels and Sinko (2006) have considered ARMA models constructed using microstructure robust measures. All of the papers mentioned above are concerned with pointwise prediction of volatility via ARMA models based on realized measures. On the other hand, there are situations in which interest may focus on predictive conditional densities, as such densities yield information not only on the conditional mean of volatility, but also on all conditional aspects of the predictive distribution. An important reason for paying attention to predictive densities of volatility is the recent development of numerous volatility-based derivative products. Examples include volatility options on various currencies such as the British pound and the Japanese Yen; and VOLAX futures, which are based upon the implied volatility of DAX index options.

The main objective of this paper is to propose a feasible, model free estimator of the conditional predictive density of integrated volatility.

From Meddahi (2003), we know that, within the context of eigenfunction stochastic volatility models, integrated volatility follows an $ARMA(p, p)$ structure, where p denotes the number of eigenfunctions. However, we only have a complete characterization of the autoregressive part of the model. Furthermore, we do not know the marginal distribution of the innovation. For these reasons, we cannot exploit the ARMA representation in order to construct predictive densities for integrated volatility. Thus, we need to follow a different route. Our approach is to construct a kernel estimator of the density of daily volatility (based on a given realized volatility measure), conditional on recent observed values of the realized measure itself. We provide general regularity conditions on the moments of the measurement error between the realized measure and integrated volatility. Given these conditions, we define a sequence of bandwidth parameters under which the kernel estimator of the conditional density is uniformly consistent. We also provide a uniform rate of convergence, which depends on the bias and variance of the kernel estimator, as well as on the measurement error. Finally, we derive the relative rate, in terms of the number of days, T, at which the bandwidth parameter and the moments of the measurement error have to approach zero, in order to ensure that all three components (bias, variance and contribution of measurement error) approach zero at the same speed. Also, we show that four well known realized measures (realized volatility; bipower variation, (Barndorff-Nielsen and Shephard, 2004, 2006); and the robust subsampled realized volatility measures of (i) Zhang et al. (2005) and (ii) Aït-Sahalia et al. (2006), Zhang (2006), Barndorff-Nielsen et al. (2008)) satisfy the conditions on the measurement error required for the uniform consistency of the estimator. This means that we can provide a feasible model free estimator of the conditional predictive density of integrated volatility even in presence of jumps or microstructure noise.

Suppose that we knew the data generating process for the instantaneous volatility. While this information suffices to characterize the autoregressive structure of the integrated volatility process, often it does not suffice to recover the "entire" data generating process. Nevertheless, in this case we can construct a kernel density estimator using the integrated volatility values simulated under the null model (and "evaluated" at the estimated parameters) instead of using a realized measure. Under mild regularity conditions, and if the null model is correct, as the sample size and the number of simulations grow at an appropriate rate, the conditional density based on kernel estimators of simulated volatility converges to the "true" conditional density of integrated volatility. A natural question is whether there is some advantage, in terms of a faster rate of convergence, in using simulated volatility rather than realized measures. We show that the answer to this question depends on the relative rate at which the number of intradaily observations, M, grows, relative to the number of days T, and on the specific realized measure used.

In order to evaluate the accuracy of our proposed estimator constructed using realized measures, we carry out a simulation experiment in which the pseudo true predictive density is compared with the one estimated using our methodology. This is done for various daily sample sizes and for a variety of different intraday data frequencies and for different data generating processes, including jumps and microstructure noise. As expected, our subsampled realized volatility measures yield substantially more accurate predictions than the other measures, when data are subject to microstructure noise. Furthermore, the predictive estimator is seen to perform quite well, overall, based on the examination of mean square error loss. We also compare the relative accuracy of predictive densities based on realized measures and on simulated integrated volatility. Finally, we provide an empirical illustration that underscores the importance of using microstructure robust measures when using data sampled at a high frequency.

The rest of the paper is organized as follows. Section 2 describes the model. Section 3 provides a uniform rate of convergence for the conditional density estimator based on a given realized measure. Section 4 provides a uniform rate of convergence for the conditional density estimator based on simulated integrated volatility, for the case in which we know the data generating process of the instantaneous volatility process. Section 5 provides conditions under which realized volatility, bipower variation and the microstructure robust measures of realized volatility satisfy the conditions on the measurement error that are required for the uniform consistency of the kernel estimator based on realized measures. Section 6 reports the results from our simulation experiment, and our empirical illustration is discussed in Section 7. Finally, Section 8 contains some concluding remarks. All proofs are gathered in the Appendix.

2. The model

The observable state variable, $Y_t = \log S_t$, where S_t denotes the price of a financial asset or the exchange rate between two currencies, is modeled as a jump diffusion process with constant drift term and variance term modeled as a measurable function of a latent factor, h_t, which is also generated by a diffusion process. Thus,

$$\mathrm{d}Y_t = m\mathrm{d}t + \mathrm{d}z_t + \sqrt{\sigma_t^2}\left(\sqrt{1-\rho^2}\mathrm{d}W_{1,t} + \rho\mathrm{d}W_{2,t}\right), \tag{1}$$

where $W_{1,t}$ and $W_{2,t}$ refer to two independent Brownian motions and volatility is modeled according to the eigenfunction stochastic volatility model of Meddahi (2001), so that

$$\sigma_t^2 = \psi(h_t) = \sum_{i=1}^{p} a_i \mathrm{P}_i(h_t)$$

$$\mathrm{d}h_t = \mu(h_t, \boldsymbol{\theta})\mathrm{d}t + \sigma(h_t, \boldsymbol{\theta})\mathrm{d}W_{2,t}, \tag{2}$$

for some $\boldsymbol{\theta} \in \boldsymbol{\Theta}$, where $\mathrm{P}_i(h_t)$ denotes the i-th eigenfunction of the infinitesimal generator $\mathcal{A}$ associated with the unobservable state variable h_t.[1] The pure jump process $\mathrm{d}z_t$ specified in (1) is such that

[1] The infinitesimal generator $\mathcal{A}$ associated with h_t is defined by

$$\mathcal{A}\phi(h_t) \equiv \mu(h_t)\phi'(h_t) + \frac{\sigma^2(h_t)}{2}\phi''(h_t)$$