



Do high-frequency measures of volatility improve forecasts of return distributions?

John M. Maheu^{a,b}, Thomas H. McCurdy^{c,d,*}

^a Department of Economics, University of Toronto, Canada

^b RCEA, Italy

^c Rotman School of Management, University of Toronto, Canada

^d CIRANO, Canada

ARTICLE INFO

Article history:

Available online 6 March 2010

Keywords:

Realized volatility

Multiperiod out-of-sample prediction

Term structure of density forecasts

Stochastic volatility

ABSTRACT

Many finance questions require the predictive distribution of returns. We propose a bivariate model of returns and realized volatility (RV), and explore which features of that time-series model contribute to superior density forecasts over horizons of 1 to 60 days out of sample. This term structure of density forecasts is used to investigate the importance of: the intraday information embodied in the daily RV estimates; the functional form for $\log(RV)$ dynamics; the timing of information availability; and the assumed distributions of both return and $\log(RV)$ innovations. We find that a joint model of returns and volatility that features two components for $\log(RV)$ provides a good fit to S&P 500 and IBM data, and is a significant improvement over an EGARCH model estimated from daily returns.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

Many finance questions require a full characterization of the distribution of returns. Examples include option pricing which uses the forecast density of the underlying spot asset, or Value-at-Risk which focuses on a quantile of the forecasted distribution. Once we move away from the simplifying assumptions of Normally-distributed returns or quadratic utility, portfolio choice also requires a full specification of the return distribution.

The purpose of this paper is to study the accuracy of forecasts of return densities produced by alternative models. Specifically, we focus on the value that high frequency measures of volatility provide in characterizing the forecast density of returns. We propose new bivariate models of returns and realized volatility and explore which features of those time-series models contribute to superior density forecasts over multiperiod horizons out of sample.

Andersen and Bollerslev (1998), Andersen et al. (2001b), Andersen et al. (2001a), Barndorff-Nielsen and Shephard (2002) and Meddahi (2002), among others,¹ have established the theoretical and empirical properties of the estimation of quadratic variation for a broad class of stochastic processes in finance. Although

theoretical advances continue to be important, part of the research in this new field has focused on the time-series properties and forecast improvements that realized volatility provides. Examples include Andersen et al. (2003), Andersen et al. (2007), Andersen et al. (2004), Ghysels and Sinko (2006), Ghysels et al. (2006), Koopman et al. (2005), Maheu and McCurdy (2002, 2007), Martens et al. (2003) and Taylor and Xu (1997).

Few papers have studied the benefits of incorporating RV into the return distribution. Andersen et al. (2003) and Giot and Laurent (2004) consider the value of RV for forecasting and for Value-at-Risk. These approaches decouple the return and volatility dynamics and assume that RV is a sufficient statistic for the conditional variance of returns. Ghysels et al. (2005) find that high frequency measures of volatility identify a risk-return tradeoff at lower frequencies. Their filtering approach to volatility measurement does not provide a law of motion for volatility and therefore multiperiod forecasts cannot be computed in that setting.

RV is an *ex post* measure of volatility and in general may not be equivalent to the conditional variance of returns. We propose bivariate models based on two alternative ways in which RV is linked to the conditional variance of returns. Since our system provides a law of motion for both return and RV at the daily frequency, multiperiod forecasts of returns and RV or the density of returns are available. The dynamics of the conditional distribution of RV will have a critical impact on the quality of the return density forecasts.

Our benchmark model is an EGARCH model of returns. This model is univariate in the sense that it is driven by one stochastic process which directs the innovations to daily returns. It does not

* Corresponding author at: Rotman School of Management, University of Toronto, Canada. Tel.: +1 416 978 3425; fax: +1 416 971 3048.

E-mail addresses: jmaheu@chass.utoronto.ca (J.M. Maheu), tmccurdy@rotman.utoronto.ca (T.H. McCurdy).

¹ Recent reviews include Andersen et al. (2009) and Barndorff-Nielsen and Shephard (2007).

allow higher-order moments of returns to be directed by a second stochastic process. Nor does it utilize any intraday information.

Two types of functional forms for the bivariate models of returns and RV are proposed. The first model uses a heterogeneous autoregressive (HAR) specification (Corsi, 2009; Andersen et al., 2007) of $\log(RV)$. A second model allows different components of $\log(RV)$ to have different decay rates (Maheu and McCurdy, 2007).

We also consider two ways to link RV to the variance of returns. First, we impose the cross-equation restriction that the conditional variance of daily returns is equal to the conditional expectation of daily RV. Second, motivated by Bollerslev et al. (2009) who model returns, bipower variation and realized jumps in a multivariate setting,² we also investigate a specification of our bivariate component model for which the variance of returns is assumed to be synonymous with RV. We label this case ‘observable stochastic volatility’ and explore whether this assumption improves the term structure of density forecasts. We also compare specifications with non-Normal versus Normal innovations for both returns and $\log(RV)$.

As in our benchmark EGARCH model, all of our bivariate models allow for so-called leverage or asymmetric effects of past negative versus positive return innovations. Our bivariate models allow for mean reversion in RV. This allows us to evaluate variance targeting for these specifications.

Our main method of model comparison uses the predictive likelihood of returns. This is the forecast density of a model evaluated at the realized return; it provides a measure of the likelihood of the data being consistent with the model. Intuitively, better forecasting models will have higher predictive likelihood values. Therefore our focus is on the relative accuracy of the models in forecasting the return density out of sample. The forecast density of the models is not available in closed form; however, we discuss accurate simulation methods that can be used to evaluate the forecast density and the predictive likelihood.

An important feature of our approach is that we can directly compare traditional volatility specifications, such as EGARCH, with our bivariate models of return and RV since we focus on a common criteria—forecast densities of returns. We generate a predictive likelihood for each out-of-sample data point and for each forecast horizon. For each forecast horizon, we can compute the average predictive likelihood where the average is computed over the fixed number of out-of-sample data points. A term structure of these average predictive likelihoods allows us to investigate the relative contributions of RV over short to long forecast horizons.

Our empirical applications to S&P 500 (Spyder) and IBM returns reveal the importance of intraday return information, the timing of information availability, and non-Normal innovations to both returns and $\log(RV)$. The main features of our results are as follows. Bivariate models that use high frequency intraday data provide a significant improvement in density forecasts relative to an EGARCH model estimated from daily data. Two-component specifications for $\log(RV)$ provide similar or better performance than HAR alternatives; both dominate the less flexible single-component version. A bivariate model of returns with Normal innovations and observable stochastic volatility directed by a 2-component, exponentially decaying function of $\log(RV)$ provides good density forecasts over a range of out-of-sample horizons for both data series. We find that adding a mixture of Normals or GARCH effects to the innovations of the $\log(RV)$ part of this specification is not statistically important for our sample of S&P 500 returns, while the addition of the mixture of Normals provides a significant improvement for IBM.

This paper is organized as follows. The next section introduces the data used to construct daily returns and daily RV. It also discusses the measurement of volatility, the adjustments to realized volatility to remove the effects of market microstructure, and a benchmark model which is based on daily return data. Our bivariate models of returns and RV, based on high-frequency intraday data, are introduced in Section 3. The calculation of density forecasts and the predictive likelihood are discussed in Section 4; results are presented in Section 5. Section 6 concludes.

2. Data and realized volatility estimation

We investigate a broadly diversified equity index (the S&P 500) and an individual stock (IBM). For the former we use the Standard & Poor’s Depository Receipt (Spyder) which is a tradable security that represents ownership in the S&P 500 Index. Since this asset is actively traded, it avoids the stale price effect associated with using the S&P 500 index at high frequencies. Transaction price data associated with both the Spyder and IBM are obtained from the New York Stock Exchange’s Trade and Quotes (TAQ) database.

Our data samples cover the period January 2, 1996 to August 29, 2007 for the Spyder and January 4, 1993 to August 29, 2007 for IBM. The shorter sample for the Spyder data was chosen based on volume of trading, for example there were many 5-minute periods with no transactions during the first years after the Spyder started trading in 1993, and a structural break in the Spyder $\log(RV)$ data in the mid 1990s (Liu and Maheu, 2008). The average number of transactions per day for the 1996–2007 sample of Spyder data was 32,971 but the volume of trades has increased substantially over the sample—especially from 2005 forward. In contrast, the average number of transactions per day for IBM shares has been more stable over our 1993–2007 sample, averaging 6011 transactions per day with a substantial increase from late 2006.

After removing errors from the transaction data,³ a 5-minute grid⁴ from 9:30 to 16:00 EST was constructed by finding the closest transaction price before or equal to each grid-point time. From this grid, 5-minute continuously compounded (log) returns were constructed. These returns were scaled by 100 and denoted as $r_{t,i}$, $i = 1, \dots, I$, where I is the number of intraday returns in day t . For our 5-minute grid, normally $I = 78$ although the market closed early on a few days. This procedure generated 228,394 5-minute returns corresponding to 2936 trading days for the S&P 500; and 286,988 5-minute returns corresponding to 3693 trading days for IBM.

The increment of quadratic variation is a natural measure of *ex post* variance over a time interval. A popular estimator is realized variance or realized volatility (RV) computed as the sum of squared returns over this time interval. The asymptotic distribution of RV has been studied by Barndorff-Nielsen and Shephard (2002) who provide conditions under which RV is an unbiased estimate.

Given the intraday returns, $r_{t,i}$, $i = 1, \dots, I$, an unadjusted daily RV estimator is

$$RV_{t,u} = \sum_{i=1}^I r_{t,i}^2. \quad (2.1)$$

However, in the presence of market-microstructure dynamics, RV

² For definition and development of bipower variation and realized jumps see, for example, Barndorff-Nielsen and Shephard (2004).

³ Data were collected with a TAQ correction indicator of 0 (regular trade) and when possible a 1 (trade later corrected), we also excluded any transaction with a sale condition of Z, which is a transaction reported on the tape out of time sequence, and with intervening trades between the trade time and the reported time on the tape. We also checked any price change that was larger than 3% and removed obvious errors.

⁴ Volatility signature plots using grids ranging from 1 min to 195 min are available on request.

Download English Version:

<https://daneshyari.com/en/article/5096764>

Download Persian Version:

<https://daneshyari.com/article/5096764>

[Daneshyari.com](https://daneshyari.com)