Available online at www.sciencedirect.com
ScienceDirect
Journal of Empirical Finance 15 (2008) 729 – 750
Journal of EMPIRICAL FINANCE
www.elsevier.com/locate/jempfin

Quantile forecasts of daily exchange rate returns from forecasts of realized volatility ☆

Michael P. Clements [a], Ana Beatriz Galvão [b], Jae H. Kim [c,*]

[a] University of Warwick, United Kingdom
[b] Queen Mary, University of London, United Kingdom
[c] Monash University, Australia

Received 30 November 2006; received in revised form 17 September 2007; accepted 10 December 2007
Available online 23 December 2007

Abstract

Quantile forecasts are central to risk management decisions because of the widespread use of Value-at-Risk. A quantile forecast is the product of two factors: the model used to forecast volatility, and the method of computing quantiles from the volatility forecasts. In this paper we calculate and evaluate quantile forecasts of the daily exchange rate returns of five currencies. The forecasting models that have been used in recent analyses of the predictability of daily realized volatility permit a comparison of the predictive power of different measures of intraday variation and intraday returns in forecasting exchange rate variability. The methods of computing quantile forecasts include making distributional assumptions for future daily returns as well as using the empirical distribution of predicted standardized returns with both rolling and recursive samples. Our main findings are that the Heterogenous Autoregressive model provides more accurate volatility and quantile forecasts for currencies which experience shifts in volatility, such as the Canadian dollar, and that the use of the empirical distribution to calculate quantiles can improve forecasts when there are shifts.

JEL classification: C32; C53; F37
Keywords: Realized volatility; Quantile forecasting; MIDAS; HAR; Exchange rates

1. Introduction

The increasing availability of high-frequency intraday data for financial variables such as stock prices and exchange rates has fuelled a rapidly growing research area in the use of realized volatility estimates to forecast daily, weekly and monthly returns volatilities and distributions. Andersen and Bollerslev (1998) showed that using realized volatility (obtained by summing the squared intraday returns) as the measure of unobserved volatility for the evaluation of daily

☆ We are grateful to Eric Ghysels and to two anonymous referees for helpful comments.
* Corresponding author. Department of Econometrics and Business Statistics, Caulfield East VIC 3145, Australia. Tel.: +61 3 99031596; fax: +61 3 99032007.

E-mail addresses: M.P.Clements@warwick.ac.uk (M.P. Clements), a.ferreira@qmul.ac.uk (A.B. Galvão), Jae.Kim@BusEco.monash.edu.au (J.H. Kim).

0927-5398/$ - see front matter
doi:10.1016/j.jempfin.2007.12.001

volatility forecasts from ARCH/GARCH models,[1] instead of the usual practice of proxying volatility using daily squared returns, suggests such forecasts are more accurate than had hitherto been found. Recent contributions have gone beyond the use of realized volatility as a measure of actual volatility for evaluation purposes, and consider the potential value of intraday returns data for forecasting volatility at lower frequencies (such as daily). Andersen et al. (2003b) set out a general framework for modelling and forecasting with high-frequency, intraday return volatilities, drawing on contributions that include Comte and Renault (1998) and Barndorff-Nielsen and Shephard (2001).[2] The (log of) the realized volatility series can be modelled using autoregressions, or vector autoregressions (VARs) when multiple related series are available. As an alternative measure to realized volatility, Barndorff-Nielsen and Shephard (2002, 2003) have proposed realized power variation – the sum of intraday absolute returns – when there are jumps in the price process. Authors such as Blair et al. (2001) have investigated adding daily realized volatility as an explanatory variable in the variance equation of GARCH models estimated on daily returns data.

Rather than modelling the aggregated intraday data (in the form of realized volatility or power variation), Ghysels et al. (2006) use the high-frequency returns directly: realized volatility is projected on to intraday squared and absolute returns using the MIDAS (MIxed Data Sampling) approach of Ghysels et al. (2004, 2006).

In the approaches exemplified by Andersen et al. (2003b) and Ghysels et al. (2004), and in a recent contribution by Koopman et al. (2005), the volatility predictions are typically compared to future realized volatilities using a loss function such as mean-squared error. The future conditional variance is taken to be quadratic variation, measured by realized volatility. Andersen et al. (2003b) justify the use of quadratic variation to measure volatility. They show that, in the absence of microstructure effects, as the sampling frequency of the intraday returns increases, the realized volatility estimates converge (almost surely) to quadratic variation. But when there are microstructure effects, the appropriate intraday sampling frequency is less clear – sampling at the highest frequencies may introduce distortions. We review issues to do with microstructure noise and investigate the appropriate sampling frequency for our exchange rate data.

Instead of comparing model forecasts as previously described, we compare models in terms of estimates of the quantiles of the distributions of future returns, such as estimates of Value-at-Risk (VaR). Our paper is closer to Giot and Laurent (2004), who compare an ARCH-type model and a model using realized volatility in terms of forecasts of Value-at-Risk. Although their particular ARCH model performs well, we narrow our study to focus exclusively on models based on realized volatility (or its constituents, intraday returns). Evidently, a quantile forecast is the product of two factors: the model used to forecast volatility, and the method of computing quantiles from the volatility forecasts. In this paper we calculate and evaluate quantile forecasts of the daily exchange rate returns of five currencies. We consider the contributions of the volatility forecasting models and the method of obtaining quantiles to the overall accuracy of the quantile forecasts. We evaluate models based on estimates of daily volatility obtained from the intraday data, and models that use the intraday data directly, along with an autoregression in realized volatility as a benchmark. These models are chosen as they have been used in recent analyses of the predictability of daily realized volatility to good effect, although there are many other models that could have been included: see for example the models in Giot and Laurent (2004). Our aim is to focus on the factors that appear to give good high-frequency quantile forecasts of exchange rates. For this purpose, a small number of volatility forecasting models will suffice.

We will assess in addition the implications of different ways of computing quantiles from the volatility estimates and forecasts, including making distributional assumptions about expected daily returns, as well as using the empirical distribution of predicted standardized returns using both rolling and recursive samples. We also take into account the role of updating the models' parameter estimates during the out-of-sample period as a way of countering potential breaks in the volatility process, and the impact this has on the quantile forecasts. Our main findings are that the Heterogenous Autoregressive (HAR) model (see Corsi (2004)) provides more accurate volatility and quantile forecasts for currencies which experience shifts in volatility, such as the Canadian dollar, and that the use of the empirical distribution to calculate quantiles can improve forecasts when there are shifts.

The plan of the remainder of the paper is as follows. The next section briefly reviews intraday-based volatility measures, and the data. Section 3 discusses the leading volatility forecasting models in the recent literature, and Section 4 the computation and evaluation of quantile forecasts. Section 5 presents the empirical results, and Section 6 some concluding remarks.

[1] See Engle (1982), Bollerslev (1986), and Bollerslev et al. (1994).

[2] Related contributions include: Andersen et al. (2001, 2003a), with applications to exchange rates; Barndorff-Nielsen and Shephard (2002, 2003), on asymptotic theory and inference. See Poon and Granger (2003) for a recent review.