Available online at www.sciencedirect.com
ScienceDirect
International Journal of Forecasting 23 (2007) 307–320
www.elsevier.com/locate/ijforecast

Forecasting realized exchange rate volatility by decomposition

Markku Lanne *

Department of Economics, P.O. Box 17 (Arkadiankatu 7), FI-00014 University of Helsinki, RUESG and HECER, Finland

Abstract

We compare forecasts of the realized volatility of the exchange rate returns of the Euro against the U.S. Dollar and the Japanese Yen obtained both directly and through decomposition. Decomposing the realized volatility into its continuous sample path and jump components, and modeling and forecasting them separately instead of directly forecasting the realized volatility, is shown to lead to improved out-of-sample forecasts. Moreover, the gains in forecast accuracy are fairly robust with respect to the details of the decomposition, but the jump component should probably not be defined too tightly.

Keywords: Realized volatility; Mixture of distributions; Aggregation; Jumps; Exchange rates

1. Introduction

Recently, Andersen, Bollerslev, and Diebold (in press) suggested modeling and forecasting the realized volatility of exchange rate, stock and bond returns by extracting the component due to jumps and including it as an explanatory variable in the heterogenous autoregressive (HAR) regression model of Müller et al. (1997) and Corsi (2003). In some cases, the jump component turned out to be highly significant, and considerable increases in the coefficient of determination were observed. This suggests that gains in forecasting the realized volatility could be made by separately modeling and forecasting the jump and continuous sample path components, and obtaining forecasts of the realized volatility as their sum instead of considering the aggregate realized volatility, as conjectured by Andersen et al. (in press). The purpose of this paper is to study whether such an approach really would be beneficial and whether the potential gains in forecast accuracy depend on the way the decomposition is carried out. To this end, we examine the returns of the Euro against the U.S. Dollar and Japanese Yen (EUR/USD and EUR/YEN returns henceforth). To model the realized volatility and the continuous components, we use the mixture multiplicative error model (mixture-MEM henceforth) previously shown by Lanne (2006) to fit well to comparable exchange rate data. One benefit of this model is that it cannot produce negative forecasts, unlike the HAR regression model used by Andersen et al. (in press). The jump components are modeled by means of the standard Markov-switching model. While the results for the two exchange rate volatility series are remarkably similar, there are also some notable differences.

The potential improvement in forecast accuracy due to decomposition can be seen to result from two factors.

* Tel.: +358 9 191 28731; fax: +358 9 191 28736.
E-mail address: markku.lanne@helsinki.fi.

0169-2070/$ - see front matter
doi:10.1016/j.ijforecast.2007.02.001

First, once the variation due to jumps has been eliminated from the realized volatility series, the process of the remaining continuous sample path component may be more easily captured, i.e., its process may be more easily estimable. Second, the jump component itself may contain predictable variation that contributes toward the forecast of the realized volatility. We show that, at least with these data, statistically significant gains in out-of-sample forecast accuracy can be made by decomposing, and this finding is fairly robust with respect to the details of the decomposition. However, if the jump component is very tightly defined, i.e., it takes nonzero values on only the days with the very greatest jumps, it has very little predictable variation so that virtually all the gains in forecast accuracy come from the improvements in estimating the process of the continuous component, causing somewhat less accurate forecasts at the one-day horizon. This finding extends the results of Andersen et al. (in press), which defined the jumps as being the very largest ones. By building a model for the jump component, we are able to forecast at longer horizons, in contrast to Andersen et al. (in press), and our results demonstrate gains in accuracy in that case as well. Although the results are clear in showing the benefits of the decomposition, the diagnostic tests suggest that as far as the jump component is concerned, even further improvements might be attainable by the use of more sophisticated models. While the results are in a sense specific to the chosen econometric models, they should be rather general, in that the mixture-MEM model has previously been shown to fit comparable exchange rate data at least as well as relevant alternatives in the literature, and diagnostic checks here also indicate its adequacy.

The plan of the paper is as follows. In Section 2 the decomposition methods put forth by Andersen et al. (in press) are reviewed and applied to the EUR/USD data. In Section 3 the mixture-MEM and Markov-switching models are introduced and the estimation results are reported, while the forecast comparisons are presented in Section 4. Finally, Section 5 concludes.

2. Decomposition of realized volatility

In this section we discuss different decompositions of the daily return variance, and introduce the data set. As a starting point for the analysis we have the realized variance of discretely sampled Δ-period returns $r_{t,\Delta} \equiv p(t) - p(t-\Delta)$,

$$\mathrm{RV}_{t+1}(\Delta) \equiv \sum_{j=1}^{1/\Delta} r_{t+j\Delta,\Delta}^2, \tag{1}$$

where $p(t)$ is the price of the asset at time point t. As Barndorff-Nielsen and Shephard (2004) have shown, the difference between this measure and the standardized realized bi-power variation,

$$\mathrm{BV}_{t+1}(\Delta) \equiv \mu_1^{-2} \sum_{j=2}^{1/\Delta} |r_{t+j\Delta,\Delta}||r_{t+(j-1)\Delta,\Delta}|$$

where $\mu_1 \equiv \sqrt{2/\pi}$, consistently estimates the component of the total return variation due to discrete jumps. Hence, it is natural to base the decomposition of $\mathrm{RV}_{t+1}(\Delta)$ on $\mathrm{RV}_{t+1}(\Delta) - \mathrm{BV}_{t+1}(\Delta)$. As this difference can also take negative values, the measure

$$J_{t+1}(\Delta) \equiv \max[\mathrm{RV}_{t+1}(\Delta) - \mathrm{BV}_{t+1}(\Delta), 0] \tag{2}$$

suggested by Barndorff-Nielsen and Shephard (2004) can be used instead to ensure the non-negativity of the jump component. The continuous sample path component $C_{t+1}(\Delta)$ simply equals $\mathrm{RV}_{t+1}(\Delta) - J_{t+1}(\Delta)$.

One potential problem with $J_{t+1}(\Delta)$ is that it typically takes positive values too frequently to be characterized as a component due to jumps. Instead, it might be desirable to identify only the significant jumps. To this end, Andersen et al. (in press) suggest employing the following test statistic

$$Z_{t+1}(\Delta) \equiv \Delta^{-1/2} \frac{[\mathrm{RV}_{t+1}(\Delta) - \mathrm{BV}_{t+1}(\Delta)]\mathrm{RV}_{t+1}(\Delta)^{-1}}{\{(\mu_1^{-4} + 2\mu_1^{-2} - 5)\max[1, \mathrm{TQ}_{t+1}(\Delta)\mathrm{BV}_{t+1}(\Delta)^{-2}]\}^{1/2}} \tag{3}$$

the distribution of which Huang and Tauchen (2005) find to be approximated by the standard normal distribution. Here, $\mathrm{TQ}_{t+1}(\Delta)$ is the standized realized tri-power quarticity,

$$\mathrm{TQ}_{t+1}(\Delta) \equiv \Delta^{-1} \mu_{4/3}^{-3} \sum_{j=3}^{1/\Delta} |r_{t+j\Delta,\Delta}|^{4/3} |r_{t+(j-1)\Delta,\Delta}|^{4/3} |r_{t+(j-2)\Delta,\Delta}|^{4/3},$$

where $\mu_{4/3} \equiv 2^{2/3}\Gamma(7/6)\Gamma(1/2)^{-1}$ and $\Gamma(\cdot)$ is the gamma function. The idea is that the jump component defined by confidence level α, $J_{t+1,\alpha}(\Delta)$, takes positive