Journal of Econometrics 160 (2011) 48–57

Contents lists available at ScienceDirect

Journal of Econometrics

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

The role of implied volatility in forecasting future realized volatility and jumps in foreign exchange, stock, and bond markets

Thomas Busch [a,d], Bent Jesper Christensen [b,d], Morten Ørregaard Nielsen [c,d,*]

[a] Danske Bank, Denmark
[b] Aarhus University, Denmark
[c] Queen's University, Canada
[d] CREATES, Denmark

ARTICLE INFO

Article history:
Available online 6 March 2010

JEL classification:
C22
C32
F31
G1

Keywords:
Bipower variation
HAR (Heterogeneous autoregressive model)
Implied volatility
Jumps
Options
Realized volatility
VecHAR
Volatility forecasting

ABSTRACT

We study the forecasting of future realized volatility in the foreign exchange, stock, and bond markets from variables in our information set, including implied volatility backed out from option prices. Realized volatility is separated into its continuous and jump components, and the heterogeneous autoregressive (HAR) model is applied with implied volatility as an additional forecasting variable. A vector HAR (VecHAR) model for the resulting simultaneous system is introduced, controlling for possible endogeneity issues. We find that implied volatility contains incremental information about future volatility in all three markets, relative to past continuous and jump components, and it is an unbiased forecast in the foreign exchange and stock markets. Out-of-sample forecasting experiments confirm that implied volatility is important in forecasting future realized volatility components in all three markets. Perhaps surprisingly, the jump component is, to some extent, predictable, and options appear calibrated to incorporate information about future jumps in all three markets.

1. Introduction

In both the theoretical and empirical finance literature, volatility is generally recognized as one of the most important determinants of risky asset values, such as exchange rates, stock and bond prices, and hence interest rates. Since any valuation procedure involves assessing the level and riskiness of future payoffs, it is particularly the forecasting of future volatility from variables in the current information set that is important for asset pricing, derivative pricing, hedging, and risk management.

A number of different variables are potentially relevant for volatility forecasting. In the present paper, we include derivative prices and investigate whether implied volatilities (IV) backed out from options on foreign currency futures, stock index futures, or Treasury bond (T-bond) futures contain incremental information when assessed against volatility forecasts based on high-frequency (5-min) current and past returns on exchange rates, stock index futures, and T-bond futures, respectively.

Andersen et al. (2003) and Andersen et al. (2004) show that simple reduced form time series models for realized volatility (RV) outperform commonly used GARCH and related stochastic volatility models in forecasting future volatility. In recent work, Barndorff-Nielsen and Shephard (2004, 2006) derive a fully nonparametric separation of the continuous sample path (C) and jump (J) components of RV. Applying this technique, Andersen et al. (2007) extend the results of Andersen et al. (2003) and Andersen et al. (2004) by using past C and J as separate regressors when forecasting volatility. They show that the two components play very different roles in forecasting, and that significant gains in performance are achieved by separating them. While C is strongly serially correlated, J is distinctly less persistent, and almost not forecastable, thus clearly indicating separate roles for C and J in volatility forecasting.

In this paper, we study high-frequency (5-min) returns to the \$/DM exchange rate, S&P 500 futures, and 30 year T-bond futures, as well as monthly prices of associated futures options. Alternative volatility measures are computed from the two separate data segments, i.e., RV and its components from high-frequency returns and IV from option prices. IV is widely perceived as a natural forecast of integrated volatility over the remaining life of the option contract under risk-neutral pricing. It is also a relevant

* Corresponding address: Department of Economics, Dunning Hall Room 307, 94 University Avenue, Queen's University, Kingston, Ontario K7L 3N6, Canada.
E-mail address: mon@econ.queensu.ca (M.Ø. Nielsen).

0304-4076/\$ – see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.jeconom.2010.03.014

forecast in a stochastic volatility setting even if volatility risk is priced, and it should get a coefficient below (above) unity in forecasting regressions in the case of a negative (positive) volatility risk premium (Bollerslev and Zhou, 2006). Since options expire at a monthly frequency, we consider the forecasting of one-month volatility measures. The issue is whether *IV* retains incremental information about future integrated volatility when assessed against realized measures (RV, C, J) from the previous month. The methodological contributions of the present paper are to use high-frequency data and recent statistical techniques for the realized measures, and to allow these to have different impacts at different frequencies, when constructing the return-based forecasts that *IV* is assessed against. These innovations ensure that *IV* is put to a harder test than in previous literature when comparing forecasting performance.

The idea of allowing different impacts at different frequencies arises since realized measures covering the entire previous month very likely are not the only relevant yardsticks. Squared returns nearly one month past may not be as informative about future volatility as squared returns that are only one or a few days old. To address this issue, we apply the heterogeneous autoregressive (HAR) model proposed by Corsi (2009) for *RV* analysis and extended by Andersen et al. (2007) to include the separate *C* and *J* components of total realized volatility ($RV = C + J$) as regressors. In the HAR framework, we include *IV* from option prices as an additional regressor, and also consider separate forecasting of both *C* and *J* individually. As an additional contribution, we introduce a vector heterogeneous autoregressive (labeled VecHAR) model for joint modeling of IV, C, and J. Since *IV* is the new variable added in our study, compared to the *RV* literature, and since it may potentially be measured with error stemming from non-synchronicity between sampled option prices and corresponding futures prices, bid-ask spreads, model error, etc., we take special care in handling this variable. The simultaneous VecHAR analysis controls for possible endogeneity issues in the forecasting equations, and allows testing interesting cross-equation restrictions.

Based on in-sample Mincer and Zarnowitz (1969) forecasting regressions, we show that *IV* contains incremental information relative to both *C* and *J* when forecasting subsequent *RV* in all three markets. Furthermore, in the foreign exchange and stock markets, *IV* is an unbiased forecast. Indeed, it completely subsumes the information content of the daily, weekly, and monthly high-frequency realized measures in the foreign exchange market. Moreover, out-of-sample forecasting evidence suggests that *IV* should be used alone when forecasting monthly *RV* in all three markets. The mean absolute out-of-sample forecast error increases if any *RV* components are included in constructing the forecast.

The results are remarkable considering that *IV* by construction should forecast volatility over the entire interval through expiration of the option, whereas our realized measures exclude the non-trading (exchange closing) intervals overnight and during weekends and holidays in the stock and bond markets. Indeed, the results most strongly favor the *IV* forecast in case of foreign currency exchange rates where there is round-the-clock trading. Squared returns over non-trading intervals could be included in *RV* for the other two markets, but with lower weight since they are more noisy. Leaving them out produces conservative results on the role of *IV* and is most natural given our focus on the separation of *C* and *J*, which in practice requires high-frequency intra-day data.

Using the HAR methodology for separate forecasting of *C* and *J*, our results show that *IV* has predictive power for each. Forecasting monthly *C* is very much like forecasting *RV* itself. The coefficient on *IV* is slightly smaller, but in-sample qualitative results on which variables to include are identical. The out-of-sample forecasting evidence suggests that *IV* again should be used alone in the foreign exchange and stock markets, but that it should be combined with realized measures in the bond market. Perhaps surprisingly, even the jump component is, to some extent, predictable, and *IV* contains incremental information about future jumps in all three markets.

The results from the VecHAR model reinforce the conclusions. In particular, when forecasting *C* in the foreign exchange market, *IV* completely subsumes the information content of all realized measures. Out-of-sample forecasting performance is about unchanged for *J* but improves for *C* in all markets by using the VecHAR model, relative to comparable univariate specifications. The VecHAR system approach allows testing cross-equation restrictions, the results of which support the finding that *IV* is a forecast of total realized volatility $RV = C + J$, indeed an unbiased forecast in the foreign exchange and stock markets.

In the previous literature, a number of authors have included *IV* in forecasting regressions, and most have found that it contains at least some incremental information, although there is mixed evidence on its unbiasedness and efficiency.[1] None of these studies has investigated whether the finding of incremental information in *IV* holds up when separating *C* and *J* computed from high-frequency returns, or when including both daily, weekly, and monthly realized measures in HAR-type specifications. An interesting alternative to using individual option prices might have been to use model-free implied volatilities as in Jiang and Tian (2005). However, Andersen and Bondarenko (2007) find that these are dominated by the simpler Black–Scholes implied volatilities in terms of forecasting power.

The remainder of the paper is laid out as follows. In the next section we briefly describe realized volatility and the nonparametric identification of its separate continuous sample path and jump components. In Section 3 we discuss the derivative pricing model. Section 4 describes our data. In Section 5 the empirical results are presented, and Section 6 concludes.

2. The econometrics of jumps

We assume that the logarithm of the asset price, $p(t)$, follows the general stochastic volatility jump diffusion model

$$dp(t) = \mu(t)dt + \sigma(t)dw(t) + \kappa(t)dq(t), \quad t \geq 0. \tag{1}$$

The mean $\mu(\cdot)$ is assumed continuous and locally bounded, the instantaneous volatility $\sigma(\cdot) > 0$ is càdlàg, and $w(\cdot)$ is the driving standard Brownian motion. The counting process $q(t)$ is normalized such that $dq(t) = 1$ corresponds to a jump at time t and $dq(t) = 0$ otherwise. Hence, $\kappa(t)$ is the random jump size at time t if $dq(t) = 1$. The intensity of the arrival process for jumps, $\lambda(t)$, is possibly time-varying, but does not allow infinite activity jump processes. Note that the leverage effect is accommodated in (1) through possible dependence between $\sigma(\cdot)$ and $w(\cdot)$, see Barndorff-Nielsen, Graversen, Jacod and Shephard (2006) and Barndorff-Nielsen, Shephard, and Winkel (2006).

The quadratic variation $[p](t)$ is defined for any semimartingale by

$$[p](t) = p\lim \sum_{j=1}^{K} (p(s_j) - p(s_{j-1}))^2, \tag{2}$$

where $0 = s_0 < s_1 < \cdots < s_K = t$ and the limit is taken for $\max_j |s_j - s_{j-1}| \to 0$ as $K \to \infty$. In the model (1), we have in wide generality

$$[p](t) = \int_0^t \sigma^2(s)ds + \sum_{j=1}^{q(t)} \kappa^2(t_j), \tag{3}$$

[1] See, e.g., Jorion (1995), Xu and Taylor (1995), Covrig and Low (2003), and Pong et al. (2004) on the foreign exchange market, Day and Lewis (1992), Canina and Figlewski (1993), Lamoureux and Lastrapes (1993), Christensen and Prabhala (1998), Fleming (1998), and Blair et al. (2001) on the stock market, and Amin and Morton (1994) on the bond market.