



The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search
<http://ageconsearch.umn.edu>
aesearch@umn.edu

Papers downloaded from AgEcon Search may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.

No endorsement of AgEcon Search or its fundraising activities by the author(s) of the following work or their employer(s) is intended or implied.

Quantile Forecasts of Daily Exchange Rate Returns from
Forecasts of Realized Volatility

Michael P. Clements, Ana Beatriz Galvão and Jae H. Kim

No 777

WARWICK ECONOMIC RESEARCH PAPERS

DEPARTMENT OF ECONOMICS

THE UNIVERSITY OF
WARWICK

Quantile Forecasts of Daily Exchange Rate Returns from Forecasts of Realized Volatility

Michael P. Clements

University of Warwick

M.P.Clements@warwick.ac.uk

Ana Beatriz Galvão*

Queen Mary, University of London

a.ferreira@qmul.ac.uk

Jae H. Kim

Monash University

Jae.Kim@BusEco.monash.edu.au

November 2006

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 HAR 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.

Key words: realized volatility, quantile forecasting, MIDAS, HAR, exchange rates.

JEL codes: C32, C53, F37

*We are grateful to Eric Ghysels for helpful comments.

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 volatility forecasts from ARCH/GARCH models¹, 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, Bollerslev, Diebold and Labys (2003) 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) and Barndorff-Nielsen and Shephard (2003) have proposed realized power variation - the sum of intraday absolute returns - when there are jumps in the price process. Authors such as Blair, Poon and Taylor (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, Santa-Clara and Valkanov (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, Santa-Clara and Valkanov (2004) and Ghysels, Sinko and Valkanov (2006).

In the approaches exemplified by Andersen et al. (2003) and Ghysels et al. (2004), and in a recent contribution by Koopman, Jungbacker and Hol (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 volatil-

¹ See Engle (1982), Bollerslev (1986), and Bollerslev, Engle and Nelson (1994).

² Related contributions include: Andersen, Bollerslev, Diebold and Labys (2000) and Andersen, Bollerslev, Diebold and Labys (2001), with applications to exchange rates; Barndorff-Nielsen and Shephard (2002) and Barndorff-Nielsen and Shephard (2003), on asymptotic theory and inference. See Poon and Granger (2003) for a recent review.

ity. Andersen et al. (2003) 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.

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. 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 HAR 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.

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 model 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.

2 Data and Volatility Measures

2.1 Exchange rate data

We have used five spot exchange rates: the Australian dollar (AU), Canadian dollar (CA), Euro (EU), U.K. pound (UK), and Japanese yen (JP), all vis-à-vis the U.S. dollar, from 4 Jan. 1999 to 31 October 2003. Following Andersen et al. (2001) and Andersen et al. (2003), 30-minute intraday returns are used to calculate the realized volatility estimates, implying that $M = 48$. This sampling frequency provides a balance between the market microstructure frictions (or noise) from high frequency sampling and the accuracy of the continuous record asymptotics from low frequency sampling. The intraday returns are calculated as the first difference of the logarithmic average of the bid-ask quotes over the 30-minute interval. Weekends, public holidays, and other inactive trading days are excluded from the sample, following Andersen et al. (2003). This gives a total of 1240 trading days, each with 47 intraday observations for a 24-hour trading day.³

2.2 Estimates of volatility

In the recent literature, volatility is often measured using realized volatility, which for daily volatility is calculated by summing up intraday squared returns:

$$RV_i = [y_M]_i^{[2]} \equiv \sum_{j=1}^M y_{j,i}^2, \quad (1)$$

where $y_{j,i}$ is the j th of M intra-day returns on day i . In the absence of microstructure effects, as M increases to infinity, the realized volatility given in (1) converges to the underlying integrated volatility, which is a natural volatility measure. Similarly, five-day realized volatility is calculated by summing squared returns over a five-day period.

$$\log RV_{t,t+5} = \log \sum_{i=1}^5 RV_{t+i}.$$

A number of studies have suggested that lags of measures of intraday variation other than realized volatility may have predictive power for realized variation. Ghysels, Santa-Clara and Valkanov (2006) and Forsberg and Ghysels (2004) propose the absolute and power variation, whilst Andersen, Bollerslev and Diebold (2005) argue for separating out a ‘jump’ component from the measure of intraday variation.

³The data source is the SIRCA (Securities Industry Research Centre of Asia-Pacific), <http://www.sirca.org.au/>.

Realized absolute variation is defined as:

$$RAV_i = \{(1/M)^{1/2}\} \sum_{j=1}^M |y_{j,i}|.$$

Forsberg and Ghysels (2004) argue for RAV as a predictor of the volatility of stock returns, on the grounds that it may be better able to capture the persistence of stock-return volatility. It can be shown that RAV is immune to jumps and the sampling error is better behaved than for RV. Notwithstanding the theoretical and empirical arguments in support of RAV as a predictor of stock-return volatility, there is no evidence on whether RAV is a useful predictor of exchange rate return volatility. We fill in the empirical evidence.

Another measure of intraday variation is bipower variation (BPV), proposed by Barndorff-Nielsen and Shephard (2003). This is defined as:

$$BPV_i = (1/M) \sum_{j=1}^{M-1} |y_{j,i}| |y_{j+1,i}|.$$

BPV has been used to separate the continuous and the jump components of *RV* (Andersen et al., 2005). The jump component can be consistently estimated by the difference between the RV and BPV:

$$\{J_M\}_i = \max\langle RV_i - BPV_i, 0 \rangle.$$

However, the jumps estimated in this way may be too small to be statistically significant. To identify statistically significant jumps, Andersen et al. (2003) suggested the use of:

$$\{Z_M\}_i = \frac{\log(RV_i) - \log(BPV_i)}{\sqrt{M^{-1}(\mu^{-4} + 2\mu^{-2} - 5)\{TQ_M\}_i(BPV_i)^{-2}}},$$

which is asymptotically distributed standard normal. In the above statistic, $\{TQ_M\}_i$ is the realized tri-power quarticity, calculated as:

$$\{TQ_M\}_i = M\mu_{4/3}^{-3} \sum_{j=1}^M |y_{j,i}|^{4/3} |y_{j+1,i}|^{4/3} |y_{j+2,i}|^{4/3},$$

where $\mu_{4/3} = 2^{2/3}\Gamma(7/6)/\Gamma(0.5)$ and $\Gamma(\cdot)$ denotes the gamma function. The significant jumps are then estimated as:

$$\{J_{M,\alpha}\}_i = I(\{Z_M\}_i > \Phi_\alpha)(RV_i - BPV_i),$$

and the continuous component as:

$$\{C_{M,\alpha}\}_i = I(\{Z_M\}_i \leq \Phi_\alpha)RV_i + I(\{Z_M\}_i > \Phi_\alpha)BPV_i,$$

where $I(\cdot)$ is the indicator function, and Φ_α denotes the critical value of the standard normal for a $(1 - \alpha)$ level test.

We estimate jump and continuous components using $\alpha = 0.95$. We find that jumps are present at around 28% of the sample, with some differences across currencies.

2.3 Summary Statistics

Figure 1 plots RV_i , its two components $\{C_{48,0.95}\}_i$ and $\{J_{48,0.95}\}_i$, and RAV_i (for $i = 1, \dots, 1240$), in standard deviation form, for Australian and Euro dollars. To conserve space, only the figures associated with these two currencies are reported. Figures for the other currencies, which can be obtained on request, show similar features. From RV_i and $\{C_{48,0.95}\}_i$, the stylized features of the conditional volatility of financial time series, documented in the ARCH literature, are evident for both currencies. The fluctuations of the volatility estimates over time are consistent with the presence of positive serial correlation, as are the jump estimates $\{J_{48,0.95}\}_i$. The estimates based on the power variation, RAV_i , are more conservative than RV_i and $\{C_{48,0.95}\}_i$ for both currencies.

Rather than modelling RV_i directly, we specify and estimate models for the log of the square root of realized volatility, $\log(RV_t^{1/2})$. The log transformation has been found to result in series which are closer to being normal (see Andersen et al. (2003)), facilitating modelling using standard autoregressions, for example. Table 1 presents some descriptive statistics for the daily and five-day realized volatility estimates. The values of skewness and kurtosis of log realized volatility are similar to those found by Andersen et al. (2003), Table II, for daily volatility, except for the UK pound which has higher negative skewness than the others. The realized volatility estimates show strong evidence of long-range dependence, as evidenced by the Ljung-Box test rejections. Visual inspection of the autocorrelation functions (not reported) show very slow declines, consistent with the observations made by Andersen et al. (2003) that the realized volatility estimates can be characterized by a long memory process.

We also report the same statistics for standardized returns - daily (five-day) returns divided by the square root of the daily (five-day) estimates of realized volatility. These match the findings for standardized returns of Andersen et al. (2000). Although we reject the null that log volatility and standardized returns are Gaussian, in most cases the departures from normality are likely to be small, and in terms of modelling log realized volatility at both daily and five-day frequencies we proceed as in the earlier studies.

Compared to earlier studies of exchange rates, we consider a greater number of series,⁴

⁴ Andersen et al. (2003) analysed the US Dollar - Deustch Mark and Dollar - Japanese Yen rates. We are not

and as will become apparent, the exchange rates exhibit different characteristics which creates variation in the performance of different models and methods across currencies.

3 Models for Volatility Forecasting

Ghysels, Santa-Clara and Valkanov (2006) and Forsberg and Ghysels (2004) evaluate the predictability of the volatility of equity returns (measured by realized volatility) over 5-day and 1-month horizons using a number of the recently proposed models. One of these models is a simple autoregression in the log of realized volatility, $\log RV_i^{1/2}$. The benchmark autoregressive model for direct calculation of h -step ahead forecasts is then:

$$\log(RV_{t,t+h}^{1/2}) = \mu + \left(\sum_{s=0}^{p-1} \psi_s L^s \right) \log(RV_{t-s-1,t-s}^{1/2}) + \varepsilon_{t,t+h}. \quad (2)$$

We consider two regression models that use alternative measures of intraday variation as explanatory variables: the Heterogenous Autoregressive model (HAR) proposed by Corsi (2004), and the MIxed Data Sampling (MIDAS) approach of Ghysels et al. (2004). The HAR model was used by Corsi to model the volatility of Swiss exchange rates, and has been extended by Andersen et al. (2005) to include jump components. These two models are discussed below.

3.1 MIDAS

The MIDAS approach uses highly parsimonious distributed lag polynomials to enable intraday data to be used to forecast daily data. The information content of the higher-frequency returns data is thus exploited in tightly parameterised models, and the problem of selecting the appropriate lag orders is in part automatically taken care of: see the references for details. Consistent estimates of the model's parameters result even though the data frequencies of the regressand and regressors differ: see Ghysels et al. (2004). The MIDAS regression to forecast the log of realized volatility using intraday squared returns has the form:

$$\log RV_{t,t+h}^{1/2} = \beta_0 + \beta_1 \log \left[B \left(L^{1/M}; \theta \right) \tilde{y}_t^2 \right]^{1/2} + \tilde{\varepsilon}_{t+h} \quad (3)$$

where $B \left(L^{1/M}; \theta \right) = \sum_{k=0}^K b(k; \theta) L^{k/M}$, $L^{k/M} \tilde{y}_{t-1}^2 = \tilde{y}_{t-k/M}^2$. Here the tilde under a variable such as y indicates that the series is at the intraday frequency. For example, when $k = 0$, $\tilde{y}_{t-k/M} = \tilde{y}_t$ refers to the last intraday return of day t , whereas y_t refers to the day t daily return. When $K > M$ intraday observations covering more than just the preceding day will

aware of forecasts of realized volatility for the Euro.

be included. In our application, the number of intraday squared returns is $M = 47$, so if $K = 235$, we use information of the past five days in forecasting, which is equivalent to $p = 5$ in equation (2). Instead of having $\{\tilde{y}_t^2\}$ on the RHS of (3), we also experiment with absolute intraday returns, $|\tilde{y}_t|$, as Forsberg and Ghysels (2004) found improvements in the predictability of stock return volatility from using absolute returns. Our work will determine which of absolute returns or squared returns are the more useful for predicting daily exchange rate volatility. We parameterise the lag polynomial $B(L^{1/M}; \theta)$ as an ‘Exponential Almon Lag’ following Ghysels, Sinko and Valkanov (2006), whereby:

$$b(k; \theta) = \frac{\exp(\theta_1 k + \theta_2 k^2)}{\sum_{k=1}^K \exp(\theta_1 k + \theta_2 k^2)}$$

In a sense the MIDAS model is more general than the autoregressive model in daily realized volatility (equation (2)). In the AR model, the implicit coefficients on all the intraday squared returns (or absolute returns) of the same day are constrained to be equal. Further, if the models were specified in terms of RV rather than $\log RV^{\frac{1}{2}}$ (and there was no log of the distributed lag on the RHS of (3)) then the MIDAS model would nest the AR. Viewed as a MIDAS model, the AR has a very specific lag polynomial structure, whereby the weights are given by a step function.

3.2 Heterogenous Autoregressive (HAR) Model

The heterogenous autoregressive model for realized volatility (HAR-RV) of Corsi (2004) and Andersen et al. (2005) specifies the current value of realized volatility as a linear function of past realized volatilities over different horizons, and can also be viewed as a restricted MIDAS model with step functions (see Forsberg and Ghysels (2004) and Ghysels, Sinko and Valkanov (2006)). The HAR-RV model can be written using the following simplifying notation. Define the normalized multi-period realized volatility as:

$$\overline{RV}_{i,i+s}^{1/2} = s^{-1}(RV_{i+1}^{1/2} + \dots + RV_{i+s}^{1/2}),$$

so that $s = 5$ and $s = 22$ are the weekly and monthly realized volatilities, respectively. Then the daily HAR-RV model that incorporates weekly and monthly realized volatility (in logarithmic form) can be written as:

$$\log(RV_{t,t+h}^{1/2}) = \beta_0 + \beta_D \log(\overline{RV}_t^{1/2}) + \beta_W \log(\overline{RV}_{t-5,t}^{1/2}) + \beta_M \log(\overline{RV}_{t-22,t}^{1/2}) + \epsilon_{t+h}.$$

Ignoring logs, it is clear that the coefficient on the intraday squared returns during the previous day is equal to $\beta_D + \beta_W + \beta_M$, on the intraday returns during days $t-4$ to $t-1$ is $\beta_W + \beta_M$, and

during days $t-21$ to $t-5$ is β_M . Assuming that $\beta_D, \beta_W, \beta_M > 0$, this corresponds to a MIDAS model in which the lag coefficients decline as a step function. However, it would be infeasible using an unrestricted MIDAS regression to allow for the monthly effect that is parameterised in the HAR-AR by the variable $\log(\bar{RV}_{t-22,t}^{1/2})$. Consequently, a potential advantage of the HAR-AR, or step-function MIDAS model, is that it is better able to capture long-range serial dependence in volatility.

The HAR-AR model can be extended to include jump components calculated using the notion of bipower variation of Barndorff-Nielsen and Shephard (2004). This gives the HAR-RV model, written as:

$$\log(RV_{t,t+h}^{1/2}) = \beta_0 + \beta_D \log(\bar{RV}_t^{1/2}) + \beta_W \log(\bar{RV}_{t-5,t}^{1/2}) + \beta_M \log(\bar{RV}_{t-22,t}^{1/2}) + \beta_J \log(1 + J_t^*) + \epsilon_{t+h},$$

where $J_i^* \equiv \{J_M\}_i$. Andersen et al. (2005) found the β_J coefficient to be statistically significant in most of their empirical examples. In addition to adding the jump component as above, the explanatory variables of the HAR-RV model can be decomposed into continuous and jump components. To simplify the notation again, let $C_i \equiv \{C_{M,\alpha}\}_i$ and $J_i \equiv \{J_{M,\alpha}\}_i$. The normalized multi-period jump and continuous components of realized volatility are respectively written as $C_{i,i+s} = s^{-1}(C_{i+1} + \dots + C_{i+s})$, and $J_{i,i+s} = h^{-1}(J_{i+1} + \dots + J_{i+s})$. Utilising the multi-period jump components separately gives the daily HAR-RV-CJ model of Andersen et al. (2005), written (in logarithmic form) as:

$$\begin{aligned} \log(RV_{t,t+h}) = & \beta_0 + \beta_{CD} \log(C_t) + \beta_{CW} \log(C_{t-5,t}) + \beta_{CM} \log(C_{t-22,t}) \\ & + \beta_{JC} \log(1 + J_t) + \beta_{JW} \log(1 + J_{t-5,t}) + \beta_{JM} \log(1 + J_{t-22,t}) + \epsilon_{t+h}. \end{aligned}$$

Andersen et al. (2005) find that most of the jump component coefficients in the HAR-RV-CJ model are statistically insignificant, and that the continuous components provide most of the predictability of the model.

The HAR can easily be specified for absolute returns, e.g.:

$$\log(RV_{t,t+h}^{1/2}) = \beta_0 + \beta_D \log(\bar{RAV}_t) + \beta_W \log(\bar{RAV}_{t-5,t}) + \beta_M \log(\bar{RAV}_{t-22,t}) + \epsilon_{t+h},$$

where $\bar{RAV}_{t-s,t}$ is the normalized multi-period absolute variation.

4 Methods for Computing and Evaluating Quantile forecasts

The models in the previous section deliver forecasts of log daily volatility over the next h days. Following Forsberg and Ghysels (2004), we obtain predicted volatility using the approximation:

$$\widehat{RV}_{t,t+h}^{1/2} = \exp \left(\widehat{\log(RV)}_{t,t+h}^{1/2} \right).$$

Conditional quantiles $q_{t,t+h}$ can be obtained by ‘inverting’ the distribution function $F_t(y) = \Pr(y_{t,t+h} \leq y | \mathcal{F}_t)$, where $y_{t,t+h}$ is the sum of daily exchange rate returns from day $t+1$ to $t+h$, and \mathcal{F}_t is the information set at t . They are computed for a given probability α so that $F_t(q_{t,t+h}) = \alpha$. Assuming that the returns are unpredictable, we have the following process for the returns $y_{t,t+h} = \varepsilon_{t,t+h}$, where $\varepsilon_{t,t+h} = \widehat{RV}_{t,t+h}^{1/2} z_{t+h}$ and z_{t+h} is *iid*. The predicted α -quantile is:

$$\hat{q}_{t,t+h} = \widehat{RV}_{t,t+h}^{1/2} F_t^{-1}(\alpha).$$

Therefore, the predicted quantiles are based on the predicted volatility but they also depend on the assumption on the predictive density $F_t(y_{t,t+h})$.

4.1 Methods for Computing the Predictive Density

The simplest method to compute $F_t^{-1}(\alpha)$ is to assume a distribution for the daily returns. Table 1 presented descriptive statistics of the standardized return $y_{t,t+h}/RV_{t,t+h}^{1/2}$, and suggests that a standard normal distribution may be a reasonable approximation. In this case, we can assume that daily returns are $N(0, RV_{t,t+h})$, so that z_t is standard normal, we have that $F_t = \Phi$. Then the quantiles with probabilities α is:

$$\left\{ z_\alpha \widehat{RV}_{t,t+h}^{1/2}, \right\} \quad (4)$$

where $z_\gamma = \Phi^{-1}(\gamma)$. The assumptions of Gaussianity of the predictive density and the unpredictability of returns underlie the popular Riskmetrics model of J.P. Morgan (1995), where $\widehat{RV}_{t,t+h}^{1/2}$ is computed recursively by a EWMA.

The assumption of normality could be replaced by a Student t assumption, or any other parametric distribution. See Bao, Lee and Saltoğlu (2004) for a discussion of some of the possibilities. In this paper, we also use a Student t with 8 degrees of freedom to capture fatter tails than the normal, although there is no strong evidence of this characteristic in the statistics of Table 1, at least for the full sample.

If standardized returns are reasonably well approximated by a normal distribution, then setting $F_t = \Phi$ should mean that improvements in volatility forecasting accuracy are associated with quantile coverage rates closer to nominal levels. That is, there is a close association between good volatility forecasts, and good quantile forecasts. If the specific distributional assumption that is adopted is poor, quantile forecasts may be improved by using instead the empirical distribution function (EDF) of the standardized returns. If the EDF is used, then it seems likely that the association between the performance of the volatility and derived quantile forecasts may be looser, in the sense that the quantile forecasts of models with relatively inaccurate

volatility forecasts may not be much worse than the quantiles from models with more accurate volatility forecasts.

Granger, White and Kamstra (1989a) suggest calculating quantiles from \widehat{Q} , the EDF of the standardized returns, $y_{t,t+h}/\widehat{RV}_{t,t+h}^{1/2}$, such that the α quantile is given by:

$$\left\{ \widehat{Q}^{-1}(\alpha) \widehat{RV}_{t,t+h}^{1/2} \right\} \quad (5)$$

Here, $\widehat{Q}^{-1}(\gamma)$ is the γ -quantile of the EDF of the standardized returns, assuming that daily returns are unpredictable in mean.

We calculate EDFs in two ways: using recursive and rolling samples of previous forecasts. To see what this means, assume that the complete sample is divided into T in-sample and n out-of-sample observations. The predicted quantiles $\hat{q}_{t,t+h}$ are computed for $t = T, T+1, \dots, T+n-h$, giving $n - (h-1)$ forecasts of length h . The EDF \widehat{Q}_t employed to compute $\hat{q}_{t,t+h}$ uses rT observations of the standardized returns $y_{t,t+h}/\widehat{RV}_{t,t+h}^{1/2}$ where $r \in (0, 1)$. In our empirical exercise, we have $r = 0.23$ implying that we use 200 observations. These observations are obtained using h -step ahead forecasts of volatility $\widehat{RV}_{t,t+h}^{1/2}$ from $t = rT+1, \dots, T-h$, assuming that the model was estimated on the the sample up to T . The difference between the rolling and recursive schemes for the computation of \widehat{Q}_t is the inclusion of the past observations of the standardized returns $y_{t+h}/\widehat{RV}_{t,t+h}^{1/2}$ while computing $\hat{q}_{t,t+h}$: the rolling scheme (q_{rol}) uses moving windows of size rT and the recursive (q_{rec}) always increases the sample adding the new observation of the standardized return at each forecast origin.

4.2 Evaluating predicted quantiles

However obtained, the quantile forecasts are evaluated by comparing their actual coverage against their nominal coverage rates. The actual coverage rates are given by $C_{\alpha,h} = E[1(y_{t,t+h} < q_{t,t+h})]$, which are estimated by $\hat{C}_{\alpha,h} = \frac{1}{n} \sum_{t=1}^n 1(y_{t,t+h} < \hat{q}_{t,t+h})$, where $t = 1, \dots, n$ indexes the forecasts. Correct unconditional coverage can be tested by a simple likelihood ratio test of whether $\hat{C}_{\alpha,h}$ is significantly different from the nominal proportion α : see e.g., Christoffersen (1998), Granger et al. (1989a). In this paper we evaluate the accuracy of VaR forecasts using the ‘tick’ or check function. The expected loss of an h -step ahead forecast made by forecaster m is defined as:

$$L_{\alpha,h,m} = E \left[\alpha - 1(y_{t,t+h} < q_{t,t+h}^m(\alpha)) \right] [y_{t,t+h} - q_{t,t+h}^m(\alpha)] \quad (6)$$

which is estimated by:

$$\hat{L}_{\alpha,h,m} = \frac{1}{n} \sum_{t=1}^n [\alpha - 1(y_{t,t+h} < \hat{q}_{t,t+h}^m(\alpha))] [y_{t,t+h} - \hat{q}_{t,t+h}^m(\alpha)].$$

The weight given to the difference between the observed return and forecasted quantile is $1 - \alpha$ when the observed return is lower than the α -quantile, but only α when the observed return exceeds the quantile. This loss function is the basis for the conditional quantile encompassing test of Giacomini and Komunjer (2005). Their test compares two rival sets of quantile forecasts to see whether either encompasses the other. However, preliminary results (unreported) indicate that the differences in the loss function for quantile forecasts from different models/methods across time are typically small given the size of the out-of-sample period that we use in this paper ($n = 340$). Another difficulty is that their test is based on the numerical estimation of parameters inside discontinuous moment conditions, which is numerically challenging given that the differences are small. Instead our forecast evaluation focuses on an unconditional test of equal forecast accuracy.

In order to see whether differences in the value of (6) across different sets of VaR forecasts are significantly different from each other, we use the testing procedure of Diebold and Mariano (1995) to make pairwise comparisons⁵ between sets of VaR forecasts using the tick loss function. The loss differential is defined as:

$$d_{t,\alpha,h} = [\alpha - 1(y_t < \hat{q}_{t,t+h}^a(\alpha))] [y_t - \hat{q}_{t,t+h}^a(\alpha)] - [\alpha - 1(y_t < \hat{q}_{t,t+h}^b(\alpha))] [y_t - \hat{q}_{t,t+h}^b(\alpha)].$$

The null that forecaster a is as accurate as forecaster b can be tested using:

$$\frac{\bar{d}_{\alpha,h}}{\sqrt{var(\bar{d}_{\alpha,h})}} \sim N(0, 1),$$

where $\bar{d}_{\alpha,h}$ is the average over t of $d_{t,\alpha,h}$. Under the alternative, we specify a one-sided test, so that rejection of the null indicates that forecaster b is more accurate than forecaster a . For $h > 1$ we use the Newey-West estimator for the variance, and a truncation lag of $h - 1$. By allotting only a relatively small fraction of our total observations to the forecast period, we are able to side-step issues related to the effects of in-sample parameter estimation uncertainty on the distribution of the test statistic (see West (2006) for a discussion).

⁵If we had a larger set of rival forecasts, it would be sensible to use the reality-check approach of White (2000). As it is, pairwise comparisons of the small set of rival forecasts enables us to more clearly see which features of the data help explain the relative forecast performances of the models.

5 Empirical Results

The objective of this empirical section is to observe which forecasting models of realized volatility and methods for computing quantile forecasts are more accurate, and to relate these findings to the underlying properties of the exchange rate series. In the first section, we focus on forecasting the volatility of exchange rate returns. In the second section we consider volatility and quantile forecasting and the potential benefits of updating the parameters of the forecasting models over the out-of-sample period. The third section evaluates the different methods of computing quantile forecasts for a given volatility forecasting model, and the fourth compares these results to those for the AR model. The final section relates the results to the properties of the individual exchange rates.

The available sample is divided into two, so that the out-of-sample period is around 1/4 of the total sample (a bit more than a year). Similar divisions into in and out-of-sample observation periods are made by Andersen et al. (2003) and Ghysels, Santa-Clara and Valkanov (2006).

5.1 Comparing Volatility Forecasts with Fixed Forecasting scheme

In this section we present both an in and out-of-sample comparison of the accuracy of volatility forecasts using the models and predictors discussed in section 3. Table 2 presents the in-sample R^2 and out-of-sample root mean squared forecast errors (RMSFE) for daily and weekly forecast horizons ($h = 1, 5$). Results are presented for an AR(5), MIDAS and HAR models. We compute forecasts from MIDAS regressions using squared and absolute returns. For the HAR, we use the basic specification, as well as specifications with separate continuous and jump components, and we also use RAV as the predictor. Average estimates over the currencies are also recorded.

The HAR is the best forecasting model overall, with more accurate forecasts on RMSE for AU, CA and UK. For the five-day volatility forecasts, gains of almost 30% can be found in comparison with the AR(5). The ability of HAR to capture the long-lag effects in a simple way is a likely reason for this success. The MIDAS forecasts are also better than the AR(5) for CA and the UK for the five-day horizon. Because the MIDAS(RV) and the AR(5) are based on the same information set, we conclude that there are gains to allowing the estimation procedure to choose how to aggregate the intraday data, rather than enforcing the daily averaging implicit in the AR model.

In contrast to the results of Forsberg and Ghysels (2004) for stock returns, there is no significant gain to using $\log(RAV)$ instead of the $\log(RV^{1/2})$ as the explanatory variable. Even in-sample, the average difference in R^2 between HAR using RAV instead of RV at $h = 5$ is less than 10 percentage points.

The forecast comparisons reported in this section are based on a fixed scheme - i.e., fixed coefficients in the out-of-sample period. This is standard practice in the volatility forecasting literature e.g., Giot and Laurent (2004), Andersen et al. (2003), and Ghysels, Santa-Clara and Valkanov (2006), but less so more generally. Breaks in the volatility process during the out-of-sample period, or parameter drift, may adversely affect forecast performance. Re-estimation of the models' parameters during the out-of-sample period may prove beneficial in these circumstances: see Clements and Hendry (2006) for a general discussion of structural breaks and forecasting. The next section considers two forms of updating.

5.2 Comparing Forecasting Models using Rolling and Recursive Samples

As we did not find significant differences from using different measures of intraday variation as explanatory variables in the forecasting models, in the following tables we present results using squared returns. Table 3 presents out-of-sample RMSFEs for the three forecasting models under fixed (as in Table 2), rolling (makes use of fixed windows of data to re-estimate the parameters over the out-of-sample period) and recursive (using increasing windows to re-estimate the models) forecasting schemes. In addition, we also compare the loss in predicting VaR at the 5% level with the tick function (eq. 6). The VaR calculations are based on the assumption that standardized returns are normally distributed.

With the exception of CA, the improvement in RMSFE accuracy of the volatility forecasts from updating the parameter estimates is relatively small at $h = 1$ for all three models. Large improvements are recorded at $h = 5$, and these are largest for the AR(5) model, and especially for CA. Differences of accuracy of forecasts between the rolling and recursive samples are virtually nonexistent. The effect of updating on the accuracy of the VaR forecasts is also small, with the exception of CA.

5.3 Predicting Quantiles with Different Distributional Assumptions

For a given volatility model, we calculate the tick loss of VaR forecasts based on different distributional assumptions. Because updating parameter estimates over the forecast period had little effect on quantile forecasts (with the exception of CA), we proceed to compare different methods of computing quantiles assuming a fixed forecasting scheme. The methods are described in section 4.1. We let $qnorm$ denote the method that assumes a normal distribution, and $qt8$ a t -distribution with 8 degrees of freedom. The other two methods use the EDF of the standardized returns to compute quantiles. $qroll$ computes the empirical quantiles using rolling samples of size 200 and $qrec$ uses an increasing window of observations from 200 up to $200 + n$.

Table 4 presents the results. The entries are the ratios of the loss using the specified distributional assumption to the loss when the predicted quantiles are computed assuming the normality of standardized returns. The results indicate that the normality assumption is adequate for weekly VaR calculations ($h = 5$), but that for the daily VaR's, alternative methods, such as the assumption that returns follow a t -distribution, or the use of the empirical distribution, are more accurate. For $h = 1$ and CA, the null that VaR forecasts calculated assuming normality are as accurate as those calculated using the EDF is rejected at the 10% level for each of the three models for $\alpha = 0.05$ and $\alpha = 0.025$. For EU there are fewer formal rejections but the value of loss tends to be lower when an alternative to the normality assumption is used.

Because these improvements in accuracy for CA and EU do not depend on the model of volatility, we conclude that the choice of distributional assumption to compute VaR may depend more on the behavior of the returns in the out-of-sample period than on the model selected for forecasting volatility.

5.4 Predicting Quantiles with Different Distributional Assumptions Relative to the AR Benchmark Model

Table 5 is similar to Table 4, in that it reports the tick loss of VaR forecasts based on different distributions, but the benchmark is now the AR(5) model (using the same distributional assumption as for the HAR or MIDAS). That is, the denominators of the ratios reported in Table 5 are always the loss for the quantiles computed using AR(5) volatility forecasts and the indicated method. This allows a cross-model comparison for a specific distributional assumption.

The results indicate that the choice of method (i.e., the distributional assumption) does not affect the choice of forecasting model. When a given model is more accurate than another, it remains so for whatever method is used to compute quantiles. Consider the $h = 1$ forecasts. The HAR model is more accurate than the AR(5) in forecasting the VaR of CA and UK, while MIDAS is more accurate for AU. At $h = 5$ the findings are more mixed, but there is some support in favour of HAR for EU and CA.

5.5 Explaining Country-Specific Differences in Forecast Performance

Our results indicate the largest differences across models and methods are for CA. Figure 2 presents the daily returns in the out-of-sample period (July 4, 2002 to October 27, 2003) and four different combinations of method/model. Using the normal distribution, we compute 5% VaR forecasts with volatility forecasts from the AR(5) and the HAR. We also use the rolling

method to compute the empirical distribution for obtaining quantiles with these same models. Figure 2 shows that after April 2003, the frequency of large negative returns increased. Before this point, there is almost no difference across models and methods in the VaR forecasts. After April 2003 it is apparent from the figure that there are marked differences between the HAR and AR model VaR forecasts. Figure 3 shows the 1-step ahead forecasts of realized volatility (and outturns) for all five currencies. From this figure, it is clear that the good performance of the HAR model VaR forecasts stems from the superior performance of the HAR volatility forecasts over this period. The HAR forecasts are better able to capture the general upturn in volatility. It is also apparent from figure 3 that currencies other than CA do not show such a clear level shift.

The reason why the HAR model volatility forecasts adapt more quickly than those of the AR to the higher level of volatility in the later part of the forecast period can be understood with a simplifying example. Suppose the true volatility process can be represented as an AR(2), and that the comparison of interest is of the AR(2) models forecasts with those of a mis-specified AR(1) model. Over the estimation sample period, the AR(2) is:

$$y_t = \phi + \alpha_1 y_{t-1} + \alpha_2 y_{t-2} + \varepsilon_t$$

and the AR(1) is $y_t = \delta + \beta y_{t-1} + u_t$. Then the population values of $[\delta, \beta]$ can be shown to be $\beta = \alpha_1 (1 - \alpha_2)^{-1}$, and $\delta = (1 - \beta) \phi (1 - \alpha_1 - \alpha_2)^{-1}$. Suppose now there is an upward shift in the unconditional mean of y_t over the forecast period to μ' (from $\mu = \phi (1 - \alpha_1 - \alpha_2)^{-1} = \delta (1 - \beta)^{-1}$). Some time after the shift has occurred, the (unconditional) forecast of the AR(2) will be: $\phi + (\alpha_1 + \alpha_2) \mu'$, versus for the AR(1): $\delta + \beta \mu'$. Provided that $\alpha_1, \alpha_2 > 0$, and the AR(2) is stationary, we can show that the AR(2) forecasts are closer to μ' than those of the AR(1). The HAR is akin to the AR(2), as it captures the longer-range dependence in volatility than the AR model (which corresponds to the AR(1) in our example). Since the sum of the AR coefficients in the HAR exceeds that of the AR (this is confirmed empirically for CA), the HAR forecasts exhibit greater adaptivity to the shift in the level of volatility.

Another result that arises from Figure 2 is that the use of the EDF (based on rolling windows), in place of an assumption of normality, to compute VaR improves the performance of both models, but especially that of the AR model (see Table 4).

In general, we find that a simple model for realized volatility (an AR(5)) and the use of the normal distribution give reasonable VaR estimates for the majority of the currencies we consider. However, when exchange rates are subject to unexpected increases in volatility, as in the case of CA, the HAR model is better able to adapt. We have provided a simple argument of how this might happen. When there are shifts, the use of the empirical distribution is better

than the normality assumption.

6 Conclusions

We have evaluated the forecast performance of a number of models that have recently been proposed to exploit the informational content of intraday data. The goal is initially to predict daily exchange rate volatility, and week-ahead exchange rate volatility. Relative to recent work, we have considered whether some of the results for stock market volatility also hold for exchange rate volatility, namely that absolute intraday returns have more predictability than squared returns. This does not appear to be the case. We also find that the method of parameterizing intraday returns implicit in the step-function MIDAS (i.e., HAR model) is generally superior to the MIDAS model which is not parameterized in this way. This appears to be due in part to the inclusion of monthly realized volatility in the former.

We then go beyond much of the recent literature to consider quantile forecasts. Quantile forecasts are the product of two factors: the model used to forecast volatility, and the method of computing quantiles from the volatility forecasts. However, the two aspects can be combined to generate a quantile forecast by either assuming a particular distributional assumption for expected future returns, or by using the volatility forecasts to obtain standardised returns from which an empirical distribution function can be estimated. One of our main findings is that a simple model for realized volatility (such as an AR(5)) combined with the assumption of a normal distribution for expected future returns yields reasonable VaR estimates for the majority of the currencies. The exception is CA, and we explain the different findings for CA in terms of a structural break in the underlying level of volatility, which appears to have been specific to CA.

From the point of view of a risk manager, the results of this paper suggest that realized volatility can be useful for computing Value-at-Risk forecasts. The combination of a simple autoregressive model for log realized volatility, together with the empirical distribution of (past) returns standardized by (past) predicted volatility, will in ‘normal times’ generate competitive Value-at-Risk forecasts with reasonable coverage rates, although there are preferred alternatives when there are structural shifts.

References

Andersen, T. G. and Bollerslev, T. (1998). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts, *International Economic Review* **39**: 885–905.

Andersen, T. G., Bollerslev, T. and Diebold, F. X. (2005). Roughing it up: Including jump components in the measurement modeling and forecasting of return volatility, *Northwestern University, Duke University and University of Pennsylvania, mimeo* .

Andersen, T. G., Bollerslev, T., Diebold, F. X. and Labys, P. (2000). Exchange rate returns standardized by realized volatility are (nearly) gaussian, *Multinational Finance Journal* **4**: 159–179.

Andersen, T. G., Bollerslev, T., Diebold, F. X. and Labys, P. (2001). The distribution of realized exchange rate volatility, *Journal of the American Statistical Association* **96**: 42–55.

Andersen, T. G., Bollerslev, T., Diebold, F. X. and Labys, P. (2003). Modeling and forecasting realized volatility, *Econometrica* **71**: 579–625.

Bao, Y., Lee, T.-H. and Saltoglu, B. (2004). A test for density forecast comparison with applications to risk management, *UC Riverside, mimeo* .

Barndorff-Nielsen, O. E. and Shephard, N. (2001). Non-gaussian Ornstein-Uhlenbeck-based models and some of their uses in financial economics (with discussion), *Journal of the Royal Statistical Society, Series B* **63**: 167–241.

Barndorff-Nielsen, O. E. and Shephard, N. (2002). Econometric analysis of realised volatility and its use in estimating stochastic volatility models, *Journal of the Royal Statistical Society, Series B* **64**: 252–280.

Barndorff-Nielsen, O. E. and Shephard, N. (2003). Realised power variation and stochastic volatility, *Bernoulli* **9**: 243–265.

Barndorff-Nielsen, O. E. and Shephard, N. (2004). Power and bipower variation with stochastic volatility and jumps, *Journal of Financial Econometrics* **2**: 1–8.

Blair, B. J., Poon, S. H. and Taylor, S. J. (2001). Forecasting S&P100 volatility: The incremental information content of implied volatilities and high frequency returns, *Journal of Econometrics* **105**: 5–26.

Bollerslev, T. (1986). Generalised autoregressive conditional heteroskedasticity, *Journal of Econometrics* **51**: 307–327.

Bollerslev, T. G., Engle, R. F. and Nelson, D. B. (1994). ARCH models, in R. F. Engle and D. McFadden (eds), *The Handbook of Econometrics, Volume 4*, North-Holland, Amsterdam, pp. 2959–3038.

Christoffersen, P. F. (1998). Evaluating interval forecasts, *International Economic Review* **39**: 841–862.

Clements, M. P. and Hendry, D. F. (2006). Forecasting with breaks, in G. Elliott, C. Granger and A. Timmermann (eds), *Handbook of Economic Forecasting, Volume 1. Handbook of Economics 24*, Elsevier, Horth-Holland, pp. 605–657.

Comte, F. and Renault, E. (1998). Long memory in continuous-time stochastic volatility models, *Mathematical Finance* **8**: 291–323.

Corsi, F. (2004). A simple long memory model of realized volatility, *University of Southern Switzerland* .

Diebold, F. X. and Mariano, R. S. (1995). Comparing predictive accuracy, *Journal of Business and Economic Statistics* **13**: 253–263. Reprinted in Mills, T. C. (ed.) (1999), *Economic Forecasting. The International Library of Critical Writings in Economics*. Cheltenham: Edward Elgar.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity, with estimates of the variance of united kingdom inflation, *Econometrica* **50**: 987–1007.

Forsberg, L. and Ghysels, E. (2004). Why do absolute returns predict volatility so well?, *University of North Carolina, (www.unc.edu/eghysels)* .

Ghysels, E., Santa-Clara, P. and Valkanov, R. (2004). The MIDAS touch: Mixed data sampling regression models, *Chapel Hill, mimeo* .

Ghysels, E., Santa-Clara, P. and Valkanov, R. (2006). Predicting volatility: Getting the most out of return data sampled at different frequencies, *Journal of Econometrics* **131**: 59–96.

Ghysels, E., Sinko, A. and Valkanov, R. (2006). MIDAS regressions: Further results and new directions, *Econometric Reviews, forthcoming* .

Giacomini, R. and Komunjer, I. (2005). Evaluation and combination of conditional quantile forecasts, *Journal of Business and Economic Statistics* **23**: 416–431.

Giot, P. and Laurent, S. (2004). Modelling daily value-at-risk using realized volatility and ARCH type models, *Journal of Empirical Finance* **11**: 379–398.

Granger, C. W. J., White, H. and Kamstra, M. (1989a). Interval forecasting: An analysis based upon ARCH-quantile estimators, *Journal of Econometrics* **40**: 87–96.

Granger, C., White, H. and Kamstra, M. (1989b). Interval forecasts: An analysis based upon ARCH-quantile estimators, *Journal of Econometrics* **40**: 87–96.

Koopman, S. J., Jungbacker, B. and Hol, E. (2005). Forecasting daily variability of the S&P 100 stock index using historical, realised and implied volatility measures, *Journal of Empirical Finance* **12**: 445–475.

Poon, S. H. and Granger, C. W. J. (2003). Forecasting volatility in financial markets: A review, *Journal of Economic Literature* **41**: 478–539.

West, K. D. (2006). Forecasting evaluation, in G. Elliott, C. Granger and A. Timmermann (eds), *Handbook of Economic Forecasting, Volume 1. Handbook of Economics 24*, Elsevier, Horth-Holland, pp. 99–134.

White, H. (2000). A reality check for data snooping, *Econometrica* **68**: 1097–1126.

Table 1: Descriptive Statistics for Daily Realized Volatility and Standardized Returns (4 Jan. 1999 to 31 October 2003)

	Mean	StDev	Skewness	Kurtosis	Q(20)	BJ(2)
$\log(RV_{t,t+1})^{(1/2)}$						
AU	-5.04	0.34	0.07	3.73	2862.1	24.14
CA	-5.61	0.34	0.08	3.36	3412.1	7.29
EU	-5.11	0.31	-0.008	3.99	1414.8	40.64
UK	-5.38	0.29	-0.37	3.89	1138.5	37.03
JP	-5.18	0.35	0.11	4.38	1715.3	70.91
$\log(RV_{t,t+5})^{(1/2)}$						
AU	-4.18	0.26	0.23	2.80	9945.0	17.89
CA	-4.76	0.25	0.25	3.18	112333.	131.72
EU	-4.25	0.22	0.22	3.39	7333.43	24.64
UK	-4.53	0.19	0.08	3.17	7045.5	2.85
JP	-4.31	0.26	0.72	3.58	6759.5	130.11
$R_{t,t+1}/(RV_{t,t+1})^{(1/2)}$						
AU	-0.06	0.89	-0.05	2.65	19.40	7.58
CA	-0.04	0.91	-0.03	2.48	15.86	16.35
EU	-0.01	0.96	-0.05	2.61	32.56	9.31
UK	-0.03	0.92	0.11	2.68	11.26	5.62
JP	-0.01	0.92	0.11	2.48	7.93	21.18
$R_{t,t+5}/(RV_{t,t+5})^{(1/2)}$						
AU	-0.11	0.88	-0.01	2.33	1512.9	29.14
CA	-0.07	0.91	0.07	2.55	1432.7	13.83
EU	-0.04	0.95	0.01	2.58	1720.6	24.34
UK	-0.06	0.95	0.07	2.61	1604.0	9.84
JP	-0.04	0.99	0.01	2.56	1565.6	1098

Note. Q(20) is the Ljung-Box test statistic for serial correlation up to 20 (Chi(20)) and BJ(2) is the statistic of the normality test (skewness =0 and kurtosis=3) for small samples. T = 1240.

Table 2: Comparing Forecasting Models: AR, MIDAS and HAR with RV, RAV and CJ as Predictors.

	R ² (T = 900)						RMSFE (T=900; n = 337)					
	AR(5)	M _(RV)	M _(RAV)	H _(RV)	H _(RAV)	H _(CJ)	AR(5)	M _(RV)	M _(RAV)	H _(RV)	H _(RAV)	H _(CJ)
h = 1												
AU	0.32	0.968	0.988	1.042	1.029	1.038	0.175	0.999	0.992	0.987	0.985	0.984
CA	0.12	0.918	0.991	1.034	1.066	1.041	0.161	1.035	0.957	0.879	0.836	0.877
EU	0.23	0.973	1.000	1.003	1.000	0.998	0.149	0.982	0.989	0.982	0.988	0.993
UK	0.20	1.000	0.976	1.032	1.003	1.042	0.114	0.989	0.991	0.992	0.992	0.991
JP	0.24	1.008	1.072	1.029	1.025	1.016	0.169	0.996	0.978	0.994	0.988	0.991
Av	0.221	0.978	1.007	1.028	1.021	1.026	0.154	1.001	0.981	0.966	0.956	0.966
h = 5												
AU	0.43	1.006	1.013	1.107	1.082	1.096	0.140	1.019	1.004	0.892	0.896	0.897
CA	0.17	0.963	0.993	1.274	1.282	1.319	0.157	0.890	0.835	0.706	0.656	0.703
EU	0.33	1.023	1.061	1.178	1.178	1.182	0.095	1.042	1.057	0.964	0.984	0.963
UK	0.38	0.990	1.003	1.027	1.020	1.036	0.086	0.872	0.882	0.850	0.860	0.850
JP	0.31	1.047	1.122	1.142	1.135	1.124	0.126	1.033	1.041	0.957	0.966	0.956
Av	0.325	1.009	1.039	1.127	1.118	1.128	0.121	0.971	0.959	0.863	0.857	0.863

Note: The entries for AR(5) are actual values. The entries for all other models are ratios over the AR(5) value. M is for MIDAS regression and H is for the Heterogeneous regression. Details are presented in section 3. Emboldened entries have ratios that indicate a difference larger than 10%. Av indicates the values computed for the average over currencies.

Table 3: Comparing RMSFE of volatility forecasting and Loss Function of VaR forecasts under different forecasting schemes

	Fixed			Rolling			Recursive		
	AR(5)	$M_{(RV)}$	$H_{(RV)}$	AR(5)	$M_{(RV)}$	$H_{(RV)}$	AR(5)	$M_{(RV)}$	$H_{(RV)}$
$h = 1$									
RMSFE									
AU 0.175 0.175 0.173 0.171 0.169 0.169 0.171 0.169 0.169 CA 0.161 0.167 0.142 0.139 0.141 0.128 0.141 0.143 0.129 EU 0.149 0.147 0.147 0.148 0.146 0.147 0.147 0.146 0.146 UK 0.114 0.112 0.113 0.111 0.110 0.110 0.111 0.111 0.110 JP 0.169 0.169 0.168 0.163 0.161 0.163 0.163 0.161 0.163 Av 0.154 0.154 0.148 0.146 0.146 0.143 0.147 0.146 0.144 Ratio 0.948 0.948 0.966 0.954 0.948 0.972									
VaR Loss Function									
AU 6.345 6.378 6.276 6.382 6.395 6.333 6.376 6.399 6.320 CA 5.457 5.084 5.542 5.123 4.945 5.182 5.143 4.958 5.193 EU 6.600 6.597 6.651 6.612 6.602 6.684 6.626 6.615 6.670 UK 5.181 5.135 5.190 5.189 5.140 5.198 5.180 5.140 5.193 JP 5.612 5.655 5.614 5.618 5.642 5.664 5.617 5.660 5.615 Av 5.839 5.770 5.854 5.785 5.745 5.812 5.788 5.754 5.798 Ratio 0.991 0.996 0.993 0.991 0.997 0.990									
$h = 5$									
RMSFE									
AU 0.140 0.143 0.125 0.123 0.123 0.115 0.122 0.122 0.115 CA 0.157 0.140 0.111 0.099 0.107 0.090 0.100 0.109 0.091 EU 0.095 0.099 0.091 0.089 0.091 0.087 0.088 0.090 0.086 UK 0.086 0.075 0.073 0.067 0.068 0.068 0.067 0.068 0.068 JP 0.126 0.130 0.121 0.111 0.109 0.110 0.112 0.111 0.111 Av 0.121 0.117 0.104 0.098 0.100 0.094 0.098 0.100 0.094 Ratio 0.810 0.961 0.903 0.810 0.961 0.903									
VaR Loss Function									
AU 11.730 12.148 11.640 11.823 12.212 11.715 11.822 12.208 11.706 CA 8.990 8.551 9.190 8.707 8.772 8.654 8.658 8.723 8.624 EU 11.585 11.631 11.762 11.621 11.691 11.790 11.589 11.656 11.744 UK 10.881 10.866 11.041 10.883 10.873 11.042 10.881 10.869 10.999 JP 11.378 11.612 11.643 11.367 11.583 11.524 11.407 11.614 11.575 Av 10.913 10.962 11.055 10.880 11.026 10.945 10.872 11.014 10.930 Ratio 0.997 1.006 0.990 0.996 1.005 0.989									

Note: Emboldened entries have ratios that indicate a difference larger than 10%. $T = 900$ and $n = 337$. The RMSFEs for the fixed forecasting scheme are the same as in Table 1. The entries are loss*1000. For the rolling scheme, the sample size is kept constant using a rolling window. For the recursive scheme, the sample size is increasing over the out-of-sample period. The rows marked 'Ratio' compare the rolling and the recursive schemes with the fixed scheme for the average over the currencies.

Table 4: Comparing Accuracy of VaR forecasts with Different Methods of Computing the Predictive Quantiles with Normal distribution as benchmark.

	AR(5)				HAR				MIDAS			
	qt8	qroll	qrec	qt8	qroll	qrec	qt8	qroll	qrec	qt8	qroll	qrec
h = 1												
$\alpha=0.05$												
AU	1.004 [.58]	1.013 [.72]	1.000 [.51]	1.001 [.53]	1.021 [.78]	0.997 [.40]	1.006 [.60]	1.011 [.71]	1.000 [.49]			
CA	0.924 [.00]	0.916 [.01]	0.946 [.00]	0.958 [.05]	0.963 [.07]	0.976 [.05]	0.922 [.00]	0.908 [.00]	0.940 [.00]			
EU	0.959 [.06]	0.969 [.16]	0.965 [.10]	0.957 [.05]	0.968 [.16]	0.962 [.08]	0.955 [.04]	0.972 [.17]	0.963 [.11]			
UK	1.006 [.61]	1.024 [.99]	1.014 [.99]	1.007 [.63]	1.016 [.94]	1.015 [1.0]	1.002 [.53]	1.020 [.96]	1.012 [.99]			
JP	1.018 [.81]	1.008 [.74]	1.008 [.89]	1.010 [.68]	1.010 [.73]	1.006 [.76]	1.023 [.88]	1.006 [.81]	1.005 [.86]			
$\alpha=0.025$												
AU	1.008 [.57]	1.022 [.75]	1.004 [.66]	0.995 [.45]	1.007 [.57]	0.999 [.49]	1.010 [.59]	1.024 [.76]	1.005 [.66]			
CA	0.904 [.03]	0.892 [.07]	0.958 [.06]	0.921 [.07]	0.950 [.18]	0.979 [.19]	0.902 [.02]	0.883 [.07]	0.950 [.05]			
EU	0.975 [.29]	0.976 [.20]	0.965 [.10]	0.960 [.20]	0.958 [.12]	0.955 [.11]	0.965 [.23]	0.964 [.13]	0.956 [.09]			
UK	1.010 [.59]	1.012 [.95]	1.005 [.83]	1.013 [.62]	1.018 [.99]	1.005 [.88]	1.003 [.52]	1.009 [.73]	1.004 [.61]			
JP	1.014 [.63]	1.020 [.71]	1.014 [.93]	0.999 [.49]	1.003 [.53]	0.994 [.37]	1.008 [.57]	1.006 [.56]	1.010 [.75]			
h = 5												
$\alpha=0.05$												
AU	1.054 [.95]	1.033 [.85]	1.024 [.83]	1.031 [.81]	1.030 [.88]	1.019 [.81]	1.060 [.97]	1.029 [.82]	1.020 [.77]			
CA	0.956 [.15]	0.997 [.46]	0.981 [.19]	1.023 [.74]	1.069 [1.0]	1.027 [.96]	0.942 [.09]	0.980 [.28]	0.970 [.12]			
EU	1.059 [.98]	1.025 [.98]	1.008 [.87]	1.045 [.93]	1.027 [.97]	1.007 [.94]	1.054 [.97]	1.038 [.99]	1.022 [.97]			
UK	0.979 [.28]	1.018 [.76]	0.994 [.37]	0.983 [.32]	1.036 [.91]	1.015 [.79]	0.970 [.22]	1.014 [.68]	0.996 [.44]			
JP	1.035 [.89]	1.130 [1.0]	1.043 [.92]	1.019 [.72]	1.090 [.97]	1.035 [.82]	1.040 [.92]	1.131 [1.0]	1.040 [.85]			
$\alpha=0.025$												
AU	1.141 [1.0]	1.029 [.68]	1.006 [.53]	1.110 [.99]	1.046 [.79]	1.007 [.56]	1.144 [1.0]	1.002 [.51]	0.986 [.41]			
CA	1.036 [.69]	1.144 [1.0]	1.036 [1.0]	1.064 [.82]	1.129 [1.0]	1.049 [.99]	1.012 [.56]	1.125 [1.0]	1.022 [.88]			
EU	1.115 [1.0]	1.061 [.91]	1.005 [.58]	1.116 [1.0]	1.049 [.93]	1.009 [.65]	1.118 [1.0]	1.059 [.92]	1.019 [.74]			
UK	1.032 [.75]	1.018 [.75]	1.014 [.70]	1.028 [.71]	1.020 [.73]	1.014 [.68]	1.025 [.70]	1.019 [.78]	1.011 [.69]			
JP	1.12 [1.0]	1.250 [1.0]	1.257 [1.0]	1.096 [.99]	1.216 [1.0]	1.209 [1.0]	1.114 [1.0]	1.225 [1.0]	1.248 [1.0]			

Note: The entries are ratios of the tick loss from using the indicated predictive density to using the normal distribution for the indicated model. The values in brackets are p-values for the null that VaR forecasts computed with normal distribution are at least as accurate as forecasts computed with the indicated predictive density. Emboldened entries indicate the null is rejected at the 10% level, implying that use of the specified method yields statistically more accurate VaRs than the normal distribution (for the given volatility forecasting model).


Table 5: Comparing Accuracy of Forecasting Models under Different Assumptions on the Predictive Density with AR(5) as benchmark.

	qnorm				qt8				qroll				qrec			
	H _(RV)		M _(RV)		H _(RV)		M _(RV)		H _(RV)		M _(RV)		H _(RV)		M _(RV)	
	H = 1															
$\alpha=0.05$																
AU	1.005	[.72]	0.989	[.09]	1.002	[.61]	0.990	[.04]	1.013	[.87]	0.987	[.07]	1.002	[.56]	0.989	[.06]
CA	0.932	[.00]	1.015	[.99]	0.967	[.05]	1.014	[.97]	0.979	[.04]	1.007	[.85]	0.961	[.01]	1.009	[.93]
EU	0.999	[.48]	1.008	[.88]	0.997	[.39]	1.003	[.73]	0.999	[.46]	1.011	[.92]	0.997	[.38]	1.006	[.90]
UK	0.991	[.09]	1.002	[.62]	0.993	[.08]	0.997	[.30]	0.984	[.06]	0.998	[.38]	0.992	[.22]	1.000	[.49]
JP	1.008	[.78]	1.000	[.52]	1.000	[.48]	1.005	[.85]	1.009	[.86]	0.999	[.44]	1.006	[.74]	0.998	[.41]
$\alpha=0.025$																
AU	1.008	[.72]	0.991	[.16]	0.995	[.35]	0.993	[.16]	0.993	[.28]	0.992	[.09]	1.004	[.62]	0.992	[.18]
CA	0.927	[.01]	1.020	[.97]	0.944	[.06]	1.018	[.93]	0.987	[.21]	1.011	[.76]	0.947	[.04]	1.012	[.87]
EU	1.005	[.60]	1.008	[.80]	0.989	[.13]	0.998	[.39]	0.986	[.26]	0.995	[.33]	0.994	[.38]	0.999	[.47]
UK	0.989	[.10]	0.997	[.38]	0.992	[.18]	0.990	[.17]	0.995	[.28]	0.994	[.34]	0.989	[.09]	0.996	[.37]
JP	1.010	[.71]	1.007	[.78]	0.995	[.37]	1.001	[.56]	0.993	[.37]	0.992	[.17]	0.990	[.29]	1.003	[.61]
H = 5																
$\alpha=0.05$																
AU	1.036	[.93]	0.992	[.32]	1.013	[.74]	0.998	[.40]	1.032	[.98]	0.988	[.18]	1.031	[.92]	0.989	[.25]
CA	0.951	[.04]	1.022	[.96]	1.018	[.82]	1.007	[.75]	1.020	[.95]	1.005	[.73]	0.996	[.42]	1.011	[.85]
EU	1.004	[.67]	1.015	[.94]	0.991	[.03]	1.010	[.97]	1.006	[.72]	1.028	[.99]	1.002	[.58]	1.029	[.98]
UK	0.999	[.45]	1.015	[.90]	1.003	[.70]	1.005	[.70]	1.016	[.91]	1.011	[.78]	1.020	[.94]	1.017	[.92]
JP	1.021	[.85]	1.023	[.92]	1.004	[.61]	1.028	[1.0]	0.984	[.02]	1.024	[.99]	1.012	[.79]	1.020	[1.0]
$\alpha=0.025$																
AU	1.035	[.82]	1.001	[.55]	1.007	[.64]	1.004	[.79]	1.052	[.91]	0.976	[.10]	1.036	[.92]	0.981	[.21]
CA	0.997	[.47]	1.003	[.59]	1.024	[.79]	0.980	[.00]	0.984	[.25]	0.987	[.09]	1.010	[.66]	0.990	[.13]
EU	0.989	[.01]	1.009	[.83]	0.989	[.02]	1.012	[.97]	0.978	[.07]	1.007	[.72]	0.993	[.16]	1.022	[.98]
UK	1.007	[.84]	1.003	[.58]	1.003	[.72]	0.997	[.36]	1.008	[.86]	1.004	[.61]	1.007	[.86]	1.000	[.50]
JP	1.017	[.74]	1.028	[.99]	0.999	[.49]	1.026	[.99]	0.989	[.26]	1.007	[.75]	0.978	[.05]	1.020	[.99]


Note: The entries are ratios of tick loss of the indicated volatility forecasting model against the AR model, when the predicted density is as indicated for both models for computing VaRs. The values in brackets are p-values for the null that VaR forecasts of the indicated model are no more accurate than forecasts of the AR(5). Emboldened entries signify the null is rejected at the 10% level.

AU


RV




RPV



C(0.95)




J(0.95)




EU


RV



RPV



C(0.95)



J(0.95)




Figure 1: Estimates of Daily Realized Volatility (std. dev) with Australian and Euro Intraday Exchange Rate Returns: Realized Quadratic Variation (RV); Realized Power Variation (RPV); Continuous and Jump Components. (Jan 4, 1999 – Oct. 31, 2003)




Figure 2: One-step-ahead Forecasts of VaRs at 5% of Canadian Dollar with AR and HAR models with Normal and Rolling Distributions




Figure 3: Realized Volatility and 1-step-ahead forecasts with AR(5), HAR and MIDAS for the five currencies.