

# Quantile forecasts of daily exchange rate returns from forecasts of realized volatility<sup>☆</sup>

Michael P. Clements<sup>a</sup>, Ana Beatriz Galvão<sup>b</sup>, Jae H. Kim<sup>c,\*</sup>

<sup>a</sup> *University of Warwick, United Kingdom*

<sup>b</sup> *Queen Mary, University of London, United Kingdom*

<sup>c</sup> *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.

© 2007 Elsevier B.V. All rights reserved.

*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

---

<sup>☆</sup> 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](mailto:M.P.Clements@warwick.ac.uk) (M.P. Clements), [a.ferreira@qmul.ac.uk](mailto:a.ferreira@qmul.ac.uk) (A.B. Galvão), [Jae.Kim@BusEco.monash.edu.au](mailto:Jae.Kim@BusEco.monash.edu.au) (J.H. Kim).

volatility forecasts from ARCH/GARCH models,<sup>1</sup> 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).<sup>2</sup> 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 Heterogeneous 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.

<sup>1</sup> See Engle (1982), Bollerslev (1986), and Bollerslev et al. (1994).

<sup>2</sup> 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.

## 2. Data and Volatility Measures

### 2.1. Exchange rate data

We use 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. We have 5-minute intraday returns calculated as the first difference of the logarithmic average of the bid-ask quotes over the 5-minute interval. Weekends, public holidays, and other inactive trading days are excluded from the sample, following Andersen et al. (2003b).<sup>3</sup> This gives a total of 1240 trading days. While some authors have used 30-minute intraday returns to calculate the realized volatility estimates, others have used 5-minute data.<sup>4</sup> Given the recent literature on the effects on noise of estimates of realized volatility, we investigate the appropriate sampling frequency for our data. We report the usual volatility signature plots in Fig. 1. The plots offer broad support for 30-minute sampling, as they appear to stabilize at around  $m=30$ . In addition, we follow the suggested way of choosing an approximate optimal sampling frequency  $M^*$  (the number of observations per day) of Bandi and Russell (2005, in press). The sampling frequency is chosen to minimize the MSE of the resulting daily realized volatility estimate,  $RV_t$ , conditional on the volatility sample path, by calculating  $M^*$  as (see also Zhang et al. (2005)):

$$M^* = \left( \frac{\bar{\hat{Q}}_i}{\hat{\alpha}} \right)^{\frac{1}{3}} \quad (1)$$

where:

$$\hat{\alpha} = \left( \frac{\sum_{i=1}^n \sum_{j=1}^M y_{j,i}^2}{nM} \right)^2$$

and:

$$\bar{\hat{Q}} = n^{-1} \sum_{i=1}^n \hat{Q}_i, \quad \hat{Q}_i = \frac{M_{15}}{3} \sum_{j=1}^{M_{15}} y_{j,i}^4.$$

The days are indexed by  $i=1, \dots, n$ , where  $y_{j,i}$  is the  $j$ th of  $M$  intraday observed returns on day  $i$ .  $y_{j,i}$  is typically contaminated with the microstructure noise in the price observations that underlie the returns series.  $\hat{Q}_i$  is the estimator of the quarticity of Barndorff-Nielsen and Shephard (2002), and  $\hat{\alpha}$  is the estimator of the squared second moment of the noise process  $((E\varepsilon^2)^2)$ , where  $\varepsilon$  is the additive error with which the observed price measures the true price). The estimator of  $(E\varepsilon^2)^2$  is based on the highest frequency at which the data are available (in our case, 5-minute data), while the subscript of 15 on  $M$  indicates  $\hat{Q}_i$  is estimated by sampling every 15-minutes ( $M_{15}=96$ ). As noted by Bandi and Russell (2005), the intuition behind (1) is clear: the larger the microstructure noise relative to the quarticity of the efficient price (as gauged by the respective estimates,  $\hat{\alpha}$  and  $\hat{Q}_i$ ), the less frequently returns should be sampled per day to avoid contaminating the RV measure with noise. The results of these calculations for the five exchange rates suggest sampling more frequently than every 30-minutes: sampling every 10 to 15 minutes might be appropriate. As a compromise, we take as our benchmark 30-minute sampling, as suggested by the volatility signature plots, but also check that some of the key findings are robust to sampling every 5-minutes, in deference to the optimal sampling frequency calculations and the studies that use 5-minute sampling.

Finally, in addition to the use of signature plots and estimating optimal sampling frequencies, there are other approaches that might be adopted to obtain RV measures when there is noise, as discussed in Andersen et al. (2006).

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

<sup>4</sup> For example, Andersen et al. (2003a,b) use 30-minute data, and Andersen et al. (2006) state that ‘within the class of linear realized volatility based forecast procedures, the use of an underlying 30-minute return horizon appears to provide a robust and reasonably efficient choice’. Studies using 5-minute data include Andersen et al. (2001) and Barndorff-Nielsen and Shephard (2004).

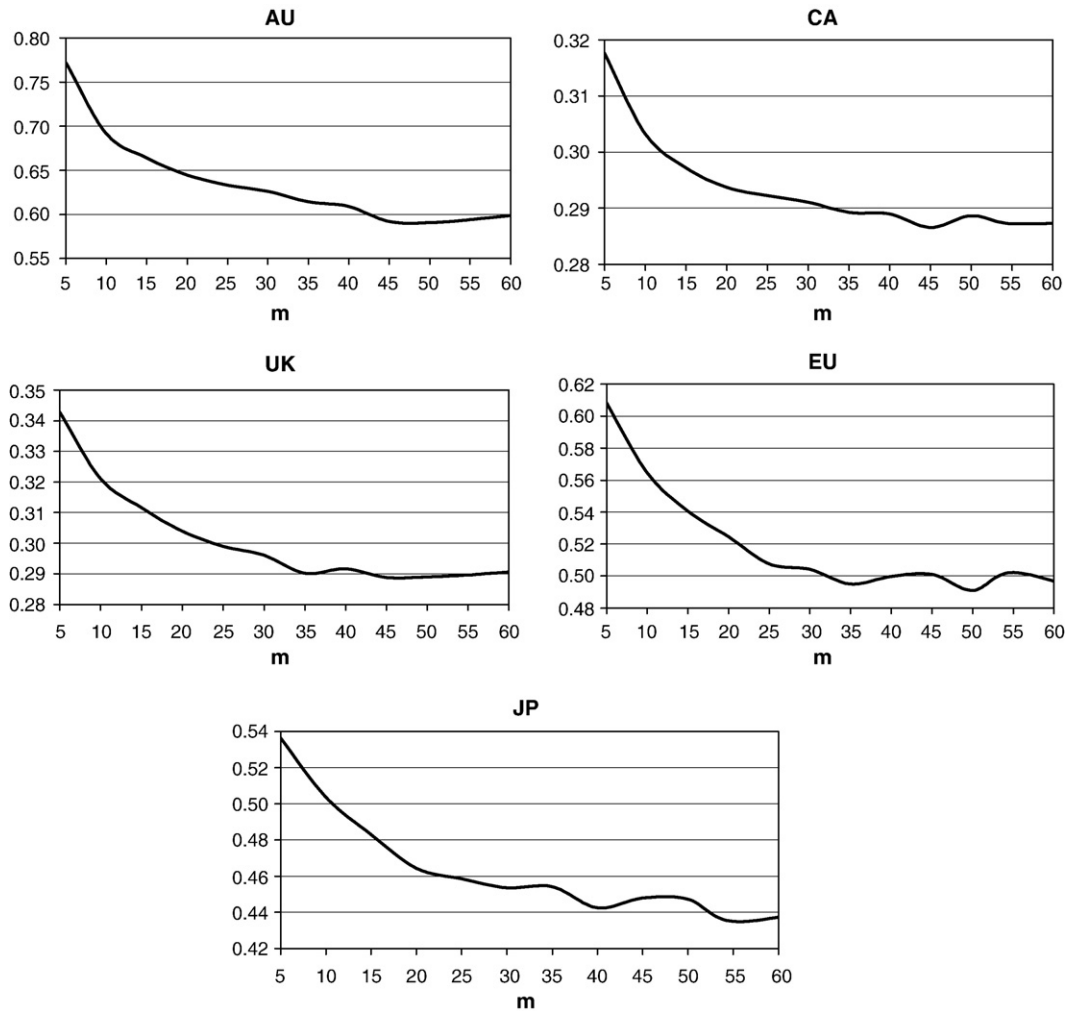


Fig. 1. Volatility signature plots for the five currencies: AU, CA, UK, EU, JP.

These include the two-scale approach of Zhang et al. (2005), as well as the Hansen and Lunde (2006) use of a Newey–West bias-correction in the presence of correlated noise. Given our main focus is not on the method of construction of the RV measure, but VaR forecasts, we do not consider these here.

## 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:

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

In the absence of microstructure effects, as  $M$  increases to infinity, the realized volatility given in Eq. (2) converges to the underlying integrated volatility, which is a natural volatility measure. As explained in the previous section, we set  $M=48$  corresponding to 30-minute sampling.

Similarly, five-day (ten-day) realized volatility is calculated by summing squared returns over a five-day (ten-day) period.

$$RV_{t,t+N_D} = \sum_{i=1}^{N_D} RV_{t+i}$$

where  $n_D=5$  (10).

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 et al. (2006) and Forsberg and Ghysels (2004) propose the absolute and power variation, whilst Andersen et al. (2005) argue for separating out a ‘jump’ component from the measure of intraday variation.

Realized absolute variation is defined as:

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

where  $\mu = \sqrt{2/\pi}$ . 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:

$$BPV_i = \mu^{-2}(1/M) \sum_{j=1}^{M-1} |y_{j,i}| |y_{j+1,i}|.$$

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. (2003b) 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  $\Phi_\alpha$  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

Fig. 2 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,  $\log(RV_i^{1/2})$ . The log transformation has been found to result in series which are closer to being normal (see Andersen et al. (2003b)), facilitating modelling using standard autoregressions, for example. Table 1 presents some descriptive statistics for the daily, five-day and ten-day realized volatility estimates. The values of skewness and kurtosis of log realized volatility are similar to those found by Andersen et al. (2003b, 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

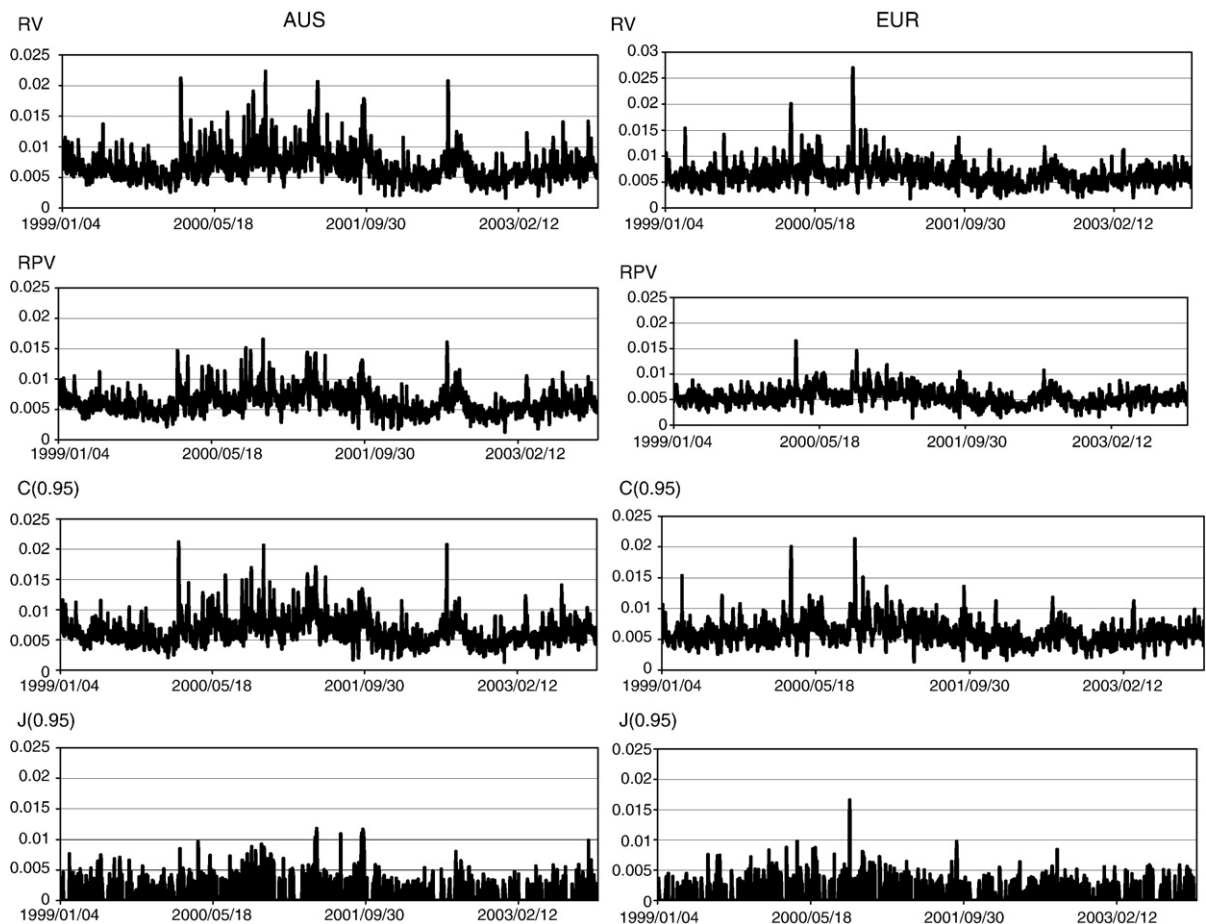


Fig. 2. Estimates of daily realized volatility (S.D.) with Australian and Euro Intraday Exchange Rate Returns: Realized Quadratic Variation (RV); Realized Power Variation (RPV); Continuous and Jump Components. (Jan 4, 1999–Oct. 31, 2003).

Table 1  
Descriptive statistics for daily realized volatility and standardized returns (4 Jan. 1999 to 31 October 2003)

	Mean	S.D.	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	11233.	131.72
EU	-4.25	0.22	0.22	3.39	7333.4	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
			$\log(RV_{t,t+10}^{(1/2)})$			
AU	-3.83	.24	.19	2.53	13959.	18.45
CA	-4.41	.23	.82	3.18	15046.	139.73
EU	-3.89	.20	.26	3.11	11495.	14.09
UK	-4.18	.17	.07	3.36	10806.	7.74
JP	-3.96	.23	.72	3.30	10234.	110.9
			$R_{t,t+1}/(RV_{t,t+1})^{(1/2)}$			
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	10.98
			$R_{t,t+10}/(RV_{t,t+10})^{(1/2)}$			
AU	-.14	.90	-0.03	2.46	3806.1	15.39
CA	-.10	.90	0.10	2.59	3510.6	10.82
EU	-.06	.99	-0.03	2.23	4197.3	30.58
UK	-.08	.98	0.12	2.59	3708.1	11.52
JP	-.07	.96	0.11	2.52	3687.9	14.39

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.

autocorrelation functions (not reported) show very slow declines, consistent with the observations made by Andersen et al. (2003b) 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 and ten-day returns divided by the square root of the relevant estimate of realized volatility. These match the findings for standardized returns of Andersen et al. (2003a). 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, five-day and ten-day frequencies we proceed as in the earlier studies.

Compared to earlier studies of exchange rates, we consider a greater number of series,<sup>5</sup> and as will become apparent, the exchange rates exhibit different characteristics which creates variation in the performance of different models and methods across currencies.

<sup>5</sup> Andersen et al. (2003b) analysed the US Dollar–Deutsch Mark and Dollar–Japanese Yen rates. We are not aware of forecasts of realized volatility for the Euro.

### 3. Models for volatility forecasting

Ghysels et al. (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,  $\log RV_t^{1/2}$ . The benchmark autoregressive model for direct calculation of h-step ahead forecasts is then:

$$\log(RV_{t,t+h}^{1/2}) = \psi_0 + \left( \sum_{s=0}^{p-1} \psi_{s+1} L^s \right) \log(RV_{t-s-1,t-s}^{1/2}) + \varepsilon_{t+h}. \quad (3)$$

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:

$$\log RV_{t,t+h}^{1/2} = \beta_0 + \beta_1 \log \left[ B(L^{1/M}; \theta) \tilde{y}_t^2 \right]^{1/2} + \tilde{\varepsilon}_{t+h} \quad (4)$$

where  $B(L^{1/M}; \theta) = \sum_{k=0}^K b(k; \theta) L^{k/M}$ ,  $L^{k/M} \tilde{y}_{t-1}^2 = \tilde{y}_{t-k/M}^2$ . Here the tilde over 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 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 Eq. (3). Instead of having  $\{\tilde{y}_t^2\}$  on the RHS of Eq. (4), 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 et al. (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 (Eq. (3)). 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  $\log RV^{1/2}$  (and there was no log of the distributed lag on the RHS of Eq. (4)) 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 the sum of a small number of past realized volatilities constructed

over different horizons, and can also be viewed as a restricted MIDAS model with step functions (see Forsberg and Ghysels, 2004 and Ghysels et al., 2006). The HAR-RV model can be written using the following simplifying notation. Define the normalized multi-period realized volatility as:

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

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:

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

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  $\beta_M$ . 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-RV by the variable  $\log \left( \overline{RV}_{t-22,t}^{1/2} \right)$ . Consequently, a potential advantage of the HAR-RV, or step-function MIDAS model, is that it is better able to capture long-range serial dependence in volatility. Corsi (2004) reports simulations that show that the HAR model is able to capture the hyperbolic decay typical of the sample autocorrelations of actual realized volatility.

The HAR-RV 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:

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

where  $J_t^* \equiv \{J_M\}_i$ . Andersen et al. (2005) found the  $\beta_J$  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  $J_{i,i+s} = s^{-1} (J_{i+1} + \dots + J_{i+s})$ . Utilising the multi-period jump components separately gives the daily HAR-RV-CJ model of Andersen et al. (2005), written (in logarithmic form) as:

$$\log \left( RV_{t,t+h} \right) = \beta_0 + \beta_{CD} \log \left( C_t \right) + \beta_{CW} \log \left( C_{t-5,t} \right) + \beta_{CM} \log \left( C_{t-22,t} \right) + \beta_{JD} \log \left( 1 + J_t \right) + \beta_{JW} \log \left( 1 + J_{t-5,t} \right) + \beta_{JM} \log \left( 1 + J_{t-22,t} \right) + \epsilon_{t+h}.$$

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.,:

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

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

Finally, we included a HAR model with the disturbance term specified as a Gaussian GARCH(1, 1) process. Augmenting the HAR with a GARCH error process gives rise to the HAR-GARCH model of Corsi et al. (2005). Those authors adopt a model of this sort after finding evidence of autoregressive conditional heteroskedasticity in the residuals of their HAR models using standard ARCH-LM tests. Rather than a standard Gaussian or Student  $t$  GARCH process, they use a standardized normal inverse Gaussian (NIG) distribution for the innovations to the GARCH process. We use the simpler formulation, but note that the findings of Corsi et al. (2005) suggest that this may be inferior to using the NIG distribution.

#### 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:

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

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  $\alpha$  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  $\varepsilon_{t,t+h} = \widehat{\text{RV}}_{t,t+h}^{1/2} z_{t+h}$  and  $z_{t+h}$  is iid. The predicted  $\alpha$ -quantile is:

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

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  $y_{t,t+h}/\widehat{\text{RV}}_{t,t+h}^{1/2}$ , and suggests that a standard normal distribution may be a reasonable approximation. In this case, we can assume that daily returns are  $N(0, \text{RV}_{t,t+h})$ ; so that  $z_t$  is standard normal, we have that  $F_t = \Phi$ . Then the quantiles with probabilities  $\alpha$  is:

$$\left\{ z_\alpha \widehat{\text{RV}}_{t,t+h}^{1/2} \right\} \quad (6)$$

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

The assumption of normality could be replaced by a Student  $t$  assumption, or any other parametric distribution. See Bao et al. (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 et al. (1989) suggest calculating quantiles from  $\widehat{Q}$ , the EDF of the standardized returns,  $y_{t,t+h}/\widehat{\text{RV}}_{t,t+h}^{1/2}$ , such that the  $\alpha$  quantile is given by:

$$\left\{ \widehat{Q}^{-1}(\alpha) \widehat{\text{RV}}_{t,t+h}^{1/2} \right\} \quad (7)$$

Here,  $\widehat{Q}^{-1}(\gamma)$  is the  $\gamma$ -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  $\hat{Q}_t$  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 sample up to  $T$ . The difference between the rolling and recursive schemes for the computation of  $\hat{Q}_t$  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 (*qroll*) uses moving windows of size  $rT$  and the recursive (*qrec*) always increases the sample adding the new observation of the standardized return at each forecast origin.

#### 4.2. Evaluating predicted quantiles

However obtained, quantile forecasts can be evaluated by comparing their actual coverage against their nominal coverage rates. The actual rates are given by  $C_{\alpha,h} = E[1(y_{t,t+h} < \hat{q}_{t,t+h})]$ , 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  $\alpha$ : see e.g., Granger et al. (1989) and Christoffersen (1998). Instead, 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:

$$L_{\alpha,h,m} = E \left[ \alpha - 1 \left( y_{t,t+h} < q_{t,t+h}^m(\alpha) \right) \right] \left[ y_{t,t+h} - q_{t,t+h}^m(\alpha) \right] \tag{8}$$

which is estimated by:

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

This is clearly related to the calculation of coverage weights, but weights the difference between the observed return and forecasted quantile by  $1-\alpha$  when the observed return is lower than the  $\alpha$ -quantile, and by  $\alpha$  when the observed return exceeds the quantile. This loss function is a natural way to evaluate quantile forecasts, as discussed by Giacomini and Komunjer (2005), who use it as the basis of a test for conditional quantile forecast encompassing. We assess whether the differences in the value of Eq. (8) across different sets of VaR forecasts are significantly different from each other, using the testing procedure of Diebold and Mariano (1995). We make pairwise comparisons<sup>6</sup> between sets of VaR forecasts. The loss differential is defined as:

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

for sets of forecasts  $a$  and  $b$ . The null that forecaster  $a$  is as accurate as forecaster  $b$  can be tested using:

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

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).

<sup>6</sup> If 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 fifth section checks the robustness of the results to the sampling frequency used to estimate the realized volatility, and last 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. (2003b) and Ghysels et al. (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, weekly and fortnightly forecast horizons ( $h=1, 5, 10$ ). Results are presented for an AR(5), MIDAS and HAR models. We compute forecasts from MIDAS regressions using squared ( $M_{(RV)}$ ) and absolute returns ( $M_{(RAV)}$ ). For the HAR, we use the basic specification ( $H_{(RV)}$ ), the one with RAV as predictor ( $H_{(RAV)}$ ), a specification with separate continuous and jump components ( $H_{(CN)}$ ), and also a HAR-GARCH ( $H_{(GA)}$ ). Average estimates over the currencies are also recorded.

The HAR is the best forecasting model overall, with more accurate forecasts on RMSE for AU and CA, and for EU and JP at  $h=5, 10$ . For the five-day volatility forecasts of CA, gains of 20% 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. From the RMSFE calculations, it is clear that the GARCH extension to the HAR does not lead to systematic improvements.

The MIDAS forecasts using RAV are better than the AR for AU and CA, where the use of RAV as the predictor outperforms using RV. The use of RAV is also better than RV in the HAR for CA. But these exceptions aside, we do not find the general improvements from using RAV reported by Forsberg and Ghysels (2004) for stock returns.

The in-sample  $R^2$ s are broadly in line with the out-of-sample RMSFEs, where the HAR is preferred, especially at  $h=5, 10$ .

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. (2003b), and Ghysels et al. (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 large differences from using different measures of intraday variation as explanatory variables taking all the currencies together, in the following tables we present results using squared returns (and forecasting models are labelled as MIDAS and HAR henceforth). We also exclude from the following tables HAR specifications that gave similar results to the basic specification (HAR-GARCH and HAR with separated continuous and jump component). 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. (8)). The VaR calculations are based on the assumption that standardized returns are normally distributed.

Table 2  
Comparing forecasting models: AR, MIDAS and HAR with RV, RAV and CJ as Predictors

	$R^2$ (T=862; common sample)							RMSFE ( $n=340$ )						
	AR	$M_{(RV)}$	$M_{(RAV)}$	$H_{(RV)}$	$H_{(RAV)}$	$H_{(CJ)}$	$H_{(GA)}$	AR	$M_{(RV)}$	$M_{(RAV)}$	$H_{(RV)}$	$H_{(RAV)}$	$H_{(CJ)}$	$H_{(GA)}$
	$h=1$													
AU	0.318	0.952	0.988	1.033	1.020	1.032	1.026	0.170	0.998	0.976	0.988	0.983	0.986	0.992
CA	0.120	0.920	0.991	1.029	1.059	1.029	1.022	0.159	1.028	0.951	0.881	0.833	0.884	0.875
EU	0.228	0.971	1.000	1.022	1.021	1.019	1.019	0.147	0.989	0.994	0.989	0.996	1.000	0.992
UK	0.204	0.987	0.976	1.029	0.997	1.045	1.023	0.111	1.002	0.994	0.993	0.992	0.998	0.989
JP	0.237	0.986	1.072	1.037	1.027	1.024	1.019	0.162	0.994	0.981	0.997	1.000	1.009	1.002
Av	0.221	0.963	1.005	1.030	1.025	1.030	1.022	0.150	1.002	0.979	0.970	0.961	0.975	0.970
	$h=5$													
AU	0.437	0.998	1.007	1.109	1.081	1.098	1.034	0.123	1.000	0.972	0.935	0.934	0.935	0.929
CA	0.167	1.005	1.047	1.312	1.325	1.350	1.113	0.136	1.018	0.941	0.791	0.725	0.789	0.841
EU	0.338	0.989	1.020	1.146	1.144	1.144	1.115	0.093	0.978	0.999	0.935	0.960	0.956	0.913
UK	0.358	1.046	1.060	1.083	1.075	1.094	0.965	0.068	0.987	0.998	0.986	1.001	0.991	1.037
JP	0.317	1.031	1.114	1.126	1.122	1.107	1.017	0.113	0.992	0.968	0.975	0.975	0.977	0.972
Av	0.323	1.014	1.050	1.155	1.149	1.158	1.049	0.107	0.995	0.976	0.924	0.919	0.930	0.938
	$h=10$													
AU	0.440	1.009	1.015	1.154	1.121	1.137	1.143	0.119	1.012	0.978	0.900	0.896	0.900	0.894
CA	0.158	0.973	1.006	1.501	1.500	1.566	1.090	0.141	1.019	0.965	0.802	0.746	0.794	0.880
EU	0.331	0.957	1.051	1.285	1.300	1.311	1.279	0.085	0.979	0.999	0.903	0.928	0.929	0.902
UK	0.358	1.004	1.032	1.133	1.131	1.143	1.066	0.061	0.992	1.006	0.991	1.004	1.001	1.030
JP	0.291	1.040	1.094	1.289	1.275	1.271	0.971	0.105	0.986	0.980	0.941	0.949	0.954	0.957
Av	0.316	0.997	1.039	1.272	1.265	1.286	1.110	0.102	0.998	0.985	0.908	0.905	0.916	0.933

Note: The entries for AR (with 5 lags) are actual values (either  $R^2$  or RMSFE). 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. (RV) means that the regressor is the realized quadratic variation. (RAV) means that the regressor is the realized absolute variation. (CJ) means that the regressors are the continuous component and jumps. (GA) is a HAR model with RV but assuming that the disturbances follow a GARCH(1,1) process. 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. The RMSFE is computed as 100 times the square root of the sum of the squared forecast errors divided by  $nh$ .

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. Larger improvements are recorded at  $h=5, 10$  on RMSFE for both AR and MIDAS, and these are again largest for CA. Differences of accuracy of forecasts between the rolling and recursive samples are virtually nonexistent. We conclude that the effect of updating is small, except for CA, and also find the effect on the accuracy of the VaR forecasts (given by VaR loss) is also small.

### 5.3. Predicting quantiles with 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. For a given volatility model, we calculate the tick loss of VaR forecasts based on different distributional assumptions. The methods are described in Section 4.1. We let  $qnorm$  denote the method that assumes a normal distribution, and  $qt\delta$  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  $200+n$ .

Tables 4a and 4b present the results for 5% and 2.5% VaRs, respectively. 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. Consider Table 4a. The results suggest that the null that the normality assumption is adequate for daily VaRs is rejected for CA using all three models. For the longer period returns, there is evidence that loss can be reduced for CA when the AR and MIDAS are used by using an assumption other than that of normality. For  $h=1$  there is an improvement in all three models' forecasts for EU when the normality assumption is abandoned,

Table 3

Comparing RMSFE of volatility forecasting and loss function of VaR forecasts under different forecasting schemes

	Fixed			Rolling			Recursive		
	AR	MIDAS	HAR	AR	MIDAS	HAR	AR	MIDAS	HAR
<i>h=1</i>									
<i>RMSFE</i>									
AU	0.170	0.169	0.168	0.175	0.174	0.173	0.175	0.173	0.173
CA	0.159	0.163	0.140	0.142	0.144	0.132	0.143	0.145	0.132
EU	0.147	0.146	0.146	0.150	0.148	0.148	0.149	0.148	0.147
UK	0.111	0.111	0.110	0.113	0.112	0.112	0.113	0.113	0.112
JP	0.162	0.161	0.162	0.168	0.167	0.167	0.169	0.168	0.168
Av	0.150	0.150	0.145	0.150	0.149	0.146	0.150	0.149	0.146
Ratio				0.999	0.994	1.009	1.000	0.996	1.009
<i>VaR loss function</i>									
AU	6.32	6.25	6.35	6.23	6.17	6.28	6.22	6.15	6.28
CA	5.43	5.55	5.07	5.02	5.13	4.97	5.05	5.14	4.97
EU	6.59	6.64	6.59	6.51	6.53	6.51	6.51	6.52	6.51
UK	5.16	5.17	5.11	5.23	5.19	5.17	5.21	5.19	5.16
JP	5.62	5.62	5.64	5.62	5.56	5.60	5.63	5.59	5.61
Av	5.82	5.84	5.75	5.72	5.72	5.71	5.72	5.72	5.71
Ratio				0.983	0.978	0.992	0.983	0.978	0.993
<i>h=5</i>									
<i>RMSFE</i>									
AU	0.123	0.123	0.115	0.128	0.128	0.119	0.126	0.127	0.118
CA	0.136	0.139	0.108	0.105	0.106	0.090	0.107	0.108	0.091
EU	0.093	0.091	0.087	0.095	0.094	0.089	0.094	0.092	0.088
UK	0.068	0.067	0.067	0.070	0.069	0.069	0.070	0.069	0.069
JP	0.113	0.112	0.110	0.114	0.114	0.113	0.117	0.116	0.114
Av	0.107	0.106	0.097	0.102	0.102	0.096	0.103	0.103	0.096
Ratio				0.962	0.961	0.984	0.965	0.965	0.986
<i>VaR loss function</i>									
AU	11.86	11.62	12.09	12.06	11.80	12.18	12.07	11.80	12.20
CA	9.08	9.19	8.58	8.66	8.66	8.77	8.63	8.62	8.74
EU	11.78	11.79	11.72	11.89	11.83	11.80	11.85	11.79	11.77
UK	10.95	10.98	10.84	10.91	10.96	10.87	10.90	10.91	10.86
JP	11.68	11.52	11.64	11.51	11.25	11.67	11.61	11.43	11.73
Av	11.07	11.02	10.97	11.01	10.90	11.06	11.01	10.91	11.06
Ratio				0.994	0.989	1.008	0.994	0.990	1.008
<i>h=10</i>									
<i>RMSFE</i>									
AU	0.119	0.121	0.107	0.122	0.124	0.110	0.120	0.121	0.109
CA	0.141	0.144	0.113	0.101	0.103	0.085	0.104	0.106	0.087
EU	0.085	0.083	0.076	0.086	0.084	0.078	0.084	0.083	0.077
UK	0.061	0.060	0.060	0.062	0.062	0.061	0.062	0.061	0.061
JP	0.105	0.104	0.099	0.103	0.102	0.099	0.107	0.105	0.101
Av	0.102	0.102	0.091	0.095	0.095	0.087	0.095	0.095	0.087
Ratio				0.926	0.927	0.951	0.933	0.932	0.956
<i>VaR loss function</i>									
AU	19.53	19.01	19.81	19.79	19.29	20.00	19.87	19.35	20.09
CA	14.98	15.13	14.52	15.01	15.03	14.77	14.87	14.90	14.68
EU	16.16	15.94	16.03	16.36	16.14	16.31	16.22	15.98	16.19
UK	13.02	13.13	13.20	12.99	13.05	13.17	13.03	13.12	13.19
JP	15.57	15.52	15.75	15.35	15.29	15.52	15.52	15.46	15.62
Av	15.85	15.75	15.86	15.90	15.76	15.95	15.90	15.76	15.95
Ratio				1.003	1.001	1.006	1.003	1.001	1.006

Table 4a

Comparing accuracy of 5% VaR forecasts with different methods of computing the predictive quantiles with normal distribution as benchmark

	AR			MIDAS			HAR											
	<i>qt8</i>	<i>qroll</i>	<i>qrec</i>	<i>qt8</i>	<i>qroll</i>	<i>qrec</i>	<i>qt8</i>	<i>qroll</i>	<i>qrec</i>									
	<i>h=1</i>																	
AU	1.006	[.60]	1.015	[.75]	1.000	[.48]	1.007	[.62]	1.011	[.71]	1.000	[.51]	1.003	[.55]	1.023	[.81]	1.001	[.52]
CA	<b>0.926</b>	<b>[.00]</b>	<b>0.917</b>	<b>[.01]</b>	<b>0.937</b>	<b>[.00]</b>	<b>0.925</b>	<b>[.00]</b>	<b>0.910</b>	<b>[.00]</b>	<b>0.932</b>	<b>[.00]</b>	<b>0.965</b>	<b>[.07]</b>	<b>0.968</b>	<b>[.10]</b>	<b>0.978</b>	<b>[.05]</b>
EU	<b>0.961</b>	<b>[.06]</b>	<b>0.970</b>	<b>[.19]</b>	0.963	[.13]	<b>0.957</b>	<b>[.05]</b>	0.970	[.16]	<b>0.963</b>	<b>[.09]</b>	<b>0.960</b>	<b>[.06]</b>	0.971	[.18]	<b>0.962</b>	<b>[.10]</b>
UK	1.007	[.63]	1.024	[.99]	1.004	[1.0]	1.003	[.55]	1.019	[.96]	1.008	[.98]	1.006	[.61]	1.018	[.96]	1.012	[.99]
JP	1.019	[.82]	1.022	[.92]	1.011	[.83]	1.026	[.92]	1.008	[.82]	1.007	[.98]	1.014	[.75]	1.015	[.77]	1.016	[.82]
	<i>h=5</i>																	
AU	1.059	[.98]	1.034	[.87]	1.027	[.83]	1.061	[.98]	1.022	[.73]	1.026	[.78]	1.037	[.86]	1.029	[.90]	1.016	[.81]
CA	0.953	[.13]	0.996	[.45]	0.985	[.25]	<b>0.946</b>	<b>[.10]</b>	0.992	[.40]	0.976	[.16]	1.024	[.75]	1.067	[1.0]	1.030	[.95]
EU	1.059	[.98]	1.030	[.99]	1.010	[.98]	1.057	[.97]	1.033	[.97]	1.010	[.91]	1.048	[.94]	1.029	[.98]	1.012	[.95]
UK	0.980	[.29]	1.029	[.88]	1.015	[.82]	0.971	[.23]	1.024	[.80]	1.005	[.59]	0.984	[.33]	1.038	[.93]	1.019	[.88]
JP	1.034	[.86]	1.110	[.99]	1.048	[.90]	1.068	[1.0]	1.088	[.99]	1.031	[.89]	1.020	[.73]	1.094	[.98]	1.045	[.86]
	<i>h=10</i>																	
AU	0.973	[.32]	1.030	[.72]	0.997	[.46]	0.971	[.32]	1.015	[.61]	1.005	[.55]	0.977	[.34]	1.043	[.73]	1.012	[.61]
CA	0.927	[.10]	0.959	[.30]	1.022	[.82]	<b>0.923</b>	<b>[.09]</b>	0.956	[.29]	1.015	[.75]	0.942	[.17]	1.006	[.54]	1.016	[.81]
EU	1.019	[.66]	1.039	[.96]	1.021	[.94]	1.018	[.65]	1.049	[.99]	1.015	[.95]	1.001	[.51]	1.042	[.95]	1.017	[.77]
UK	1.025	[.72]	1.124	[1.0]	1.036	[.91]	1.014	[.62]	1.135	[1.0]	1.033	[.87]	1.012	[.60]	1.104	[.99]	1.035	[.81]
JP	1.053	[.98]	1.166	[1.0]	1.145	[1.0]	1.055	[.98]	1.188	[1.0]	1.152	[1.0]	1.014	[.65]	1.145	[1.0]	1.120	[.99]

Note: The entries are ratios of the tick loss from using the indicated predictive density (*qt8*, *qroll* or *qrec*) to using the normal distribution (*qnorm*) 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).

although these are not always sufficiently large to result in significant test outcomes. But apart from these findings for CA and EU, we do not find statistically significant reductions in loss from using a non-normal distributional assumption for any of the other currencies using any of the three models. For the longer horizons, the loss ratios indicate that use of the EDF may be worse than using a normality assumption, a result we attribute in part to the relatively small samples available to calculate the EDFs. To a lesser degree these results hold for the 2.5% VaR (see Table 4b).

#### 5.4. Predicting quantiles with different distributional assumptions relative to the AR benchmark model

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

These results indicate that the HAR is better than the AR for CA and the UK for daily VaR irrespective of the distributional assumption, so that the choice of distribution does not affect the ranking between forecasting models in these cases. However, for CA the performance of the AR improves relative to the HAR when the EDFs are used instead of the normal distribution assumption. We also find that MIDAS is better than the AR for AU for all VaR horizons irrespective of the distributional assumption we make.

#### Notes to Table 3:

Note: Number of forecasts, *n*, is 340. The RMSFEs for the fixed forecasting scheme are the same as in Table 2. The entries are loss times 10,000. 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 headed ‘Ratio’ compare the rolling and the recursive schemes with the fixed scheme for the average over the currencies. “MIDAS” was labelled  $M_{(RV)}$  in Table 2 and “HAR” was labelled  $H_{(RV)}$ .

Table 4b

Comparing accuracy of 2.5% VaR forecasts with different methods of computing the predictive quantiles with normal distribution as benchmark

	AR			MIDAS			HAR											
	<i>qt8</i>	<i>qroll</i>	<i>qrec</i>	<i>qt8</i>	<i>qroll</i>	<i>qrec</i>	<i>qt8</i>	<i>qroll</i>	<i>qrec</i>									
	<i>h=1</i>																	
AU	1.009	[.58]	1.022	[.75]	1.012	[.81]	1.012	[.60]	1.024	[.76]	1.006	[.72]	1.001	[.51]	1.017	[.69]	1.006	[.67]
CA	<b>0.906</b>	<b>[.03]</b>	<b>0.880</b>	<b>[.06]</b>	<b>0.942</b>	<b>[.05]</b>	<b>0.905</b>	<b>[.02]</b>	<b>0.876</b>	<b>[.06]</b>	<b>0.931</b>	<b>[.04]</b>	<b>0.923</b>	<b>[.08]</b>	0.939	[.15]	0.973	[.17]
EU	0.977	[.31]	0.977	[.22]	0.968	[.16]	0.968	[.25]	0.967	[.17]	0.960	[.14]	0.965	[.23]	0.972	[.22]	0.958	[.15]
UK	1.011	[.61]	1.016	[.94]	1.007	[.76]	1.004	[.53]	1.013	[.79]	1.007	[.69]	1.015	[.64]	1.028	[.98]	1.015	[.95]
JP	1.015	[.65]	1.014	[.63]	1.011	[.65]	1.009	[.59]	1.015	[.62]	1.008	[.60]	0.999	[.49]	1.006	[.55]	0.997	[.47]
	<i>h=5</i>																	
AU	1.127	[1.0]	1.012	[.58]	0.972	[.29]	1.146	[1.0]	0.991	[.43]	0.980	[.38]	1.116	[1.0]	1.028	[.70]	1.008	[.56]
CA	1.023	[.63]	1.116	[1.0]	1.029	[.92]	1.013	[.57]	1.113	[1.0]	1.024	[.89]	1.064	[.82]	1.114	[1.0]	1.045	[1.0]
EU	1.122	[1.0]	1.030	[.73]	0.995	[.40]	1.121	[1.0]	1.027	[.72]	0.996	[.42]	1.116	[1.0]	1.041	[.89]	1.007	[.62]
UK	1.030	[.73]	1.035	[.95]	1.023	[.87]	1.028	[.72]	1.026	[.85]	1.010	[.67]	1.032	[.75]	1.025	[.78]	1.018	[.71]
JP	1.109	[1.0]	1.215	[1.0]	1.192	[1.0]	1.131	[1.0]	1.150	[1.0]	1.134	[1.0]	1.099	[1.0]	1.203	[1.0]	1.186	[1.0]
	<i>h=10</i>																	
AU	1.020	[.60]	1.040	[.72]	1.050	[.85]	1.051	[.75]	1.044	[.83]	1.050	[.85]	0.993	[.47]	1.031	[.66]	1.027	[.82]
CA	0.898	[.20]	1.081	[.81]	1.089	[.88]	0.885	[.17]	1.051	[.69]	1.086	[.87]	0.938	[.31]	1.078	[.82]	1.120	[.96]
EU	1.125	[1.0]	1.095	[.99]	1.043	[.98]	1.151	[1.0]	1.109	[.94]	1.045	[.97]	1.132	[1.0]	1.086	[.99]	1.038	[1.0]
UK	1.133	[1.0]	1.177	[.99]	1.024	[.89]	1.130	[1.0]	1.178	[.99]	1.038	[.96]	1.130	[1.0]	1.164	[.99]	1.051	[.97]
JP	1.126	[1.0]	1.225	[1.0]	1.220	[1.0]	1.138	[1.0]	1.241	[1.0]	1.262	[1.0]	1.110	[1.0]	1.226	[1.0]	1.220	[1.0]

Note: See notes to Table 4a.

### 5.5. Effect of the sampling-frequency used to calculate realized volatility

Table 6 presents a check on the robustness of some of these findings to the sampling frequency, by reporting VaR loss using RV estimates obtained using 5-minute sampling, and comparing these to the results obtained using 30-

Table 5a

Comparing accuracy of 5% VaR forecasts between forecasting models under different assumptions on the predictive density with AR as benchmark

	<i>qnorm</i>		<i>qt8</i>		<i>qrec</i>		<i>qroll</i>									
	MIDAS	HAR	MIDAS	HAR	MIDAS	HAR	MIDAS	HAR								
	<i>h=1</i>															
AU	<b>0.989</b>	<b>[.08]</b>	1.005	[.71]	<b>0.990</b>	<b>[.03]</b>	1.002	[.61]	<b>0.985</b>	<b>[.03]</b>	1.012	[.93]	<b>0.989</b>	<b>[.08]</b>	1.006	[.75]
CA	1.014	[.99]	<b>0.934</b>	<b>[.00]</b>	1.013	[.98]	<b>0.973</b>	<b>[.08]</b>	1.007	[.88]	<b>0.987</b>	<b>[.09]</b>	1.009	[.95]	<b>0.975</b>	<b>[.03]</b>
EU	1.008	[.87]	0.999	[.43]	1.003	[.72]	0.998	[.35]	1.008	[.79]	1.000	[.50]	1.008	[.79]	0.998	[.39]
UK	1.002	[.62]	<b>0.990</b>	<b>[.05]</b>	0.997	[.29]	<b>0.989</b>	<b>[.02]</b>	0.997	[.33]	<b>0.984</b>	<b>[.05]</b>	1.006	[.85]	0.998	[.43]
JP	1.006	[.80]	1.003	[.65]	1.014	[.99]	0.998	[.41]	0.992	[.23]	0.996	[.34]	1.002	[.55]	1.008	[.81]
	<i>h=5</i>															
AU	<b>0.980</b>	<b>[.06]</b>	1.019	[.83]	<b>0.982</b>	<b>[.08]</b>	0.999	[.46]	<b>0.968</b>	<b>[.01]</b>	1.014	[.77]	0.979	[.11]	1.009	[.73]
CA	1.012	[.99]	<b>0.945</b>	<b>[.04]</b>	1.004	[.81]	1.015	[.76]	1.008	[.93]	1.012	[.84]	1.003	[.76]	0.987	[.28]
EU	1.000	[.47]	0.994	[.29]	0.998	[.23]	<b>0.984</b>	<b>[.00]</b>	1.003	[.65]	0.994	[.25]	1.000	[.49]	0.997	[.36]
UK	1.003	[.60]	0.990	[.15]	0.994	[.29]	0.994	[.20]	0.998	[.41]	0.998	[.42]	0.993	[.27]	0.994	[.23]
JP	0.966	[.16]	0.997	[.37]	0.998	[.47]	<b>0.983</b>	<b>[.01]</b>	<b>0.947</b>	<b>[.00]</b>	0.982	[.08]	<b>0.950</b>	<b>[.00]</b>	0.993	[.15]
	<i>h=10</i>															
AU	<b>0.973</b>	<b>[.08]</b>	1.014	[.73]	<b>0.971</b>	<b>[.07]</b>	1.019	[.80]	<b>0.958</b>	<b>[.02]</b>	1.027	[.98]	0.981	[.11]	1.030	[.90]
CA	1.010	[.97]	0.969	[.15]	1.006	[.86]	0.985	[.33]	1.007	[.90]	1.017	[.85]	1.004	[.77]	0.964	[.12]
EU	0.986	[.11]	0.992	[.20]	0.986	[.08]	0.974	[.00]	0.997	[.39]	0.996	[.41]	0.981	[.11]	0.989	[.31]
UK	1.009	[.88]	1.014	[.81]	0.997	[.27]	1.000	[.51]	1.019	[.99]	0.996	[.43]	1.005	[.93]	1.013	[.77]
JP	0.996	[.33]	1.011	[.74]	0.998	[.39]	0.974	[.00]	1.015	[1.0]	0.993	[.26]	1.003	[.66]	0.989	[.12]

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.

Table 5b

Comparing accuracy of 2.5% VaR forecasts between forecasting models under different assumptions on the predictive density with AR as benchmark

	<i>qnorm</i>				<i>qt8</i>				<i>qrec</i>				<i>qroll</i>			
	MIDAS		HAR		MIDAS		HAR		MIDAS		HAR		MIDAS		HAR	
	<i>h=1</i>															
AU	0.991	[.15]	1.010	[.78]	0.993	[.16]	1.001	[.54]	<b>0.993</b>	<b>[.09]</b>	1.005	[.68]	<b>0.985</b>	<b>[.09]</b>	1.003	[.61]
CA	1.018	[.97]	<b>0.934</b>	<b>[.01]</b>	1.017	[.94]	<b>0.951</b>	<b>[.07]</b>	1.014	[.86]	0.997	[.40]	1.006	[.80]	<b>0.964</b>	<b>[.07]</b>
EU	1.008	[.78]	1.001	[.55]	0.998	[.36]	0.989	[.04]	0.998	[.41]	0.997	[.40]	1.000	[.51]	0.991	[.25]
UK	0.997	[.38]	<b>0.983</b>	<b>[.02]</b>	0.990	[.15]	<b>0.988</b>	<b>[.07]</b>	0.994	[.32]	0.995	[.27]	0.997	[.40]	0.991	[.11]
JP	1.022	[.99]	1.010	[.75]	1.015	[.95]	0.994	[.32]	1.023	[.99]	1.002	[.55]	1.019	[.98]	0.997	[.39]
	<i>h=5</i>															
AU	<b>0.970</b>	<b>[.09]</b>	1.003	[.56]	0.986	[.13]	0.993	[.26]	<b>0.949</b>	<b>[.02]</b>	1.019	[.89]	<b>0.978</b>	<b>[.05]</b>	1.040	[.97]
CA	1.009	[.88]	1.002	[.53]	0.999	[.46]	1.043	[.94]	1.007	[.86]	1.001	[.51]	1.004	[.73]	1.017	[.74]
EU	0.999	[.32]	<b>0.987</b>	<b>[.03]</b>	0.997	[.19]	<b>0.981</b>	<b>[.00]</b>	0.996	[.16]	0.998	[.45]	1.000	[.45]	0.999	[.47]
UK	0.990	[.31]	0.995	[.26]	0.988	[.08]	0.998	[.32]	0.981	[.20]	0.986	[.18]	0.978	[.17]	0.990	[.29]
JP	0.977	[.28]	<b>0.987</b>	<b>[.05]</b>	0.997	[.46]	<b>0.978</b>	<b>[.00]</b>	<b>0.925</b>	<b>[.00]</b>	<b>0.977</b>	<b>[.01]</b>	<b>0.929</b>	<b>[.00]</b>	<b>0.982</b>	<b>[.00]</b>
	<i>h=10</i>															
AU	<b>0.948</b>	<b>[.05]</b>	1.035	[.86]	0.976	[.13]	1.008	[.63]	0.951	[.13]	1.026	[.99]	<b>0.947</b>	<b>[.06]</b>	1.012	[.78]
CA	1.013	[.95]	0.960	[.21]	0.998	[.30]	1.003	[.52]	0.984	[.09]	0.957	[.15]	1.011	[.84]	0.988	[.31]
EU	0.976	[.10]	<b>0.969</b>	<b>[.01]</b>	0.997	[.18]	<b>0.975</b>	<b>[.00]</b>	0.988	[.12]	<b>0.962</b>	<b>[.01]</b>	0.978	[.07]	<b>0.964</b>	<b>[.09]</b>
UK	0.998	[.37]	1.000	[.49]	0.995	[.05]	0.996	[.30]	0.999	[.43]	0.989	[.22]	1.011	[.87]	1.026	[.99]
JP	0.995	[.34]	<b>0.976</b>	<b>[.02]</b>	1.007	[.91]	<b>0.962</b>	<b>[.00]</b>	1.009	[.77]	<b>0.978</b>	<b>[.02]</b>	1.029	[1.0]	<b>0.975</b>	<b>[.00]</b>

Note: See notes to Table 5a.

Table 6

Comparing Loss functions of 5% VaR forecasts using 5- and 30-min data

	Loss with 5-min data			Ratio to 30-min data		
	AR	MIDAS	HAR	AR	MIDAS	HAR
	<i>h=1</i>					
AU	7.78	7.82	7.67	1.23	1.25	1.21
CA	6.39	6.24	6.30	1.18	1.13	1.24
EU	7.26	7.23	7.19	1.10	1.09	1.09
UK	5.63	5.64	5.60	1.09	1.09	1.10
JP	5.79	5.84	5.79	1.03	1.04	1.03
Av.	6.57	6.55	6.51			
	<i>h=5</i>					
AU	19.47	19.56	19.25	1.64	1.68	1.59
CA	24.18	24.29	24.06	2.66	2.64	2.81
EU	16.65	16.59	16.48	1.41	1.41	1.41
UK	12.85	12.86	12.79	1.17	1.17	1.18
JP	12.28	12.27	12.06	1.05	1.06	1.04
Av.	17.09	17.11	16.93			
	<i>h=10</i>					
AU	27.22	27.34	26.65	1.39	1.44	1.35
CA	46.52	46.59	46.60	3.11	3.08	3.21
EU	24.14	24.13	23.84	1.49	1.51	1.49
UK	18.65	18.71	18.51	1.43	1.42	1.40
JP	16.30	16.35	16.14	1.05	1.05	1.02
Av.	26.57	26.62	26.35			

Note: The entries in the left panel are loss functions (times 10,000) of 5% VaR forecasts when the realized volatility and returns are computed with data sampled at each 5 min. Loss functions computed for data sampled at each 30 min are reported in Table 3.

Table 7

Comparing accuracy of 5% VaR forecasts between forecasting models under different assumptions on the predictive density with AR as benchmark

	<i>qnorm</i>				<i>qt8</i>				<i>qrec</i>				<i>qroll</i>			
	MIDAS		HAR		MIDAS		HAR		MIDAS		HAR		MIDAS		HAR	
	<i>h=1</i>															
AU	1.005	[.89]	<b>0.987</b>	[.00]	1.006	[.92]	<b>0.988</b>	[.00]	1.007	[.93]	1.002	[.65]	1.009	[.94]	1.002	[.63]
CA	<b>0.977</b>	[.01]	<b>0.986</b>	[.04]	<b>0.978</b>	[.03]	<b>0.987</b>	[.08]	<b>0.984</b>	[.04]	0.991	[.10]	<b>0.977</b>	[.04]	0.992	[.17]
EU	0.995	[.15]	<b>0.990</b>	[.00]	0.996	[.24]	<b>0.990</b>	[.00]	1.002	[.64]	<b>0.994</b>	[.08]	<b>0.992</b>	[.06]	0.997	[.30]
UK	1.001	[.57]	<b>0.995</b>	[.07]	1.001	[.56]	<b>0.994</b>	[.02]	0.997	[.35]	<b>0.991</b>	[.03]	1.005	[.72]	0.996	[.17]
JP	1.008	[.99]	0.999	[.44]	1.010	[1.0]	0.995	[.19]	1.010	[.94]	1.005	[.83]	0.996	[.29]	1.002	[.63]
	<i>h=5</i>															
AU	1.005	[.90]	<b>0.989</b>	[.02]	1.003	[.81]	<b>0.986</b>	[.00]	1.005	[.90]	<b>0.988</b>	[.03]	1.002	[.68]	<b>0.992</b>	[.08]
CA	1.005	[.77]	0.995	[.26]	1.006	[.84]	1.008	[.89]	1.009	[.83]	1.023	[1.0]	1.010	[.93]	1.025	[.99]
EU	0.996	[.18]	<b>0.989</b>	[.06]	0.998	[.31]	<b>0.985</b>	[.00]	0.996	[.22]	0.997	[.25]	0.999	[.38]	<b>0.990</b>	[.08]
UK	1.001	[.57]	0.995	[.12]	1.001	[.64]	<b>0.996</b>	[.07]	<b>0.995</b>	[.08]	<b>0.992</b>	[.07]	<b>0.991</b>	[.02]	<b>0.988</b>	[.00]
JP	0.999	[.36]	<b>0.982</b>	[.00]	1.004	[.96]	<b>0.979</b>	[.00]	1.001	[.60]	<b>0.984</b>	[.03]	<b>0.995</b>	[.07]	<b>0.989</b>	[.09]
	<i>h=10</i>															
AU	1.004	[.87]	<b>0.979</b>	[.00]	1.005	[.96]	<b>0.976</b>	[.00]	1.012	[.97]	1.005	[.88]	1.004	[.85]	<b>0.988</b>	[.10]
CA	1.001	[.66]	1.002	[.59]	1.004	[.81]	1.003	[.64]	1.011	[.99]	1.015	[.89]	1.000	[.47]	1.018	[.92]
EU	1.000	[.46]	0.988	[.13]	1.003	[.92]	<b>0.984</b>	[.07]	0.998	[.35]	<b>0.979</b>	[.05]	0.996	[.21]	0.997	[.43]
UK	1.003	[.83]	<b>0.992</b>	[.06]	1.002	[.70]	<b>0.989</b>	[.01]	1.001	[.58]	<b>0.985</b>	[.10]	1.015	[.99]	0.997	[.39]
JP	1.003	[.78]	0.990	[.17]	1.007	[.99]	<b>0.970</b>	[.00]	1.006	[.94]	<b>0.977</b>	[.02]	0.997	[.28]	0.995	[.30]

Note: This Table reproduces results of Table 5a with data sampled at each 5 min instead of 30 min.

minute sampling to estimate RV. The VaRs are calculated assuming normality and a fixed forecasting scheme. The MIDAS models (Eq. (4)) are estimated with  $K=1435$ , which is the usual 5 days of past data when  $M=287$ . The table shows that loss is always smaller using the 30-minute returns to calculate RV, and that the ranking of the models is generally unaffected by the switch to 5-minute returns (the ratios across models for a given currency are generally similar).

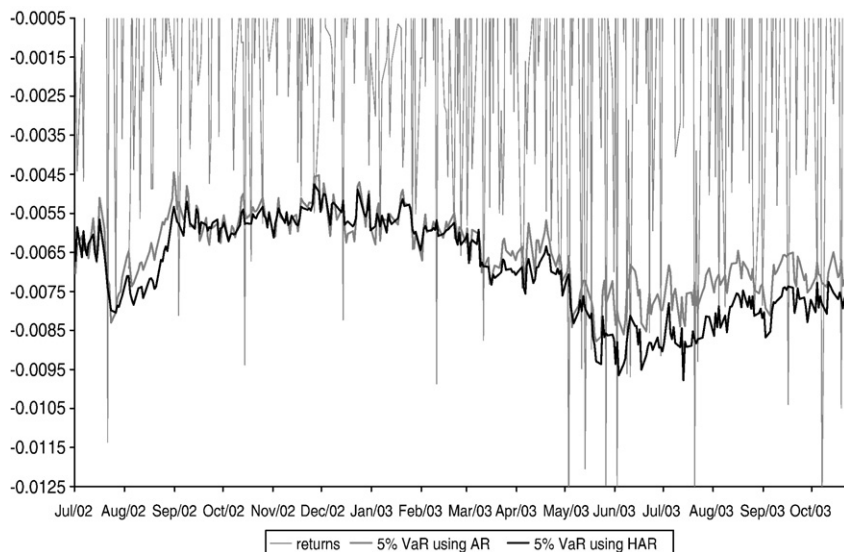


Fig. 3. One-step-ahead forecasts of VaRs at 5% of Canadian Dollar with AR and HAR models with normal distribution.




Fig. 4. Realized volatility and 1-step-ahead forecasts with AR, HAR and MIDAS for the five currencies.

We next calculate VaRs for different distributional assumptions for 5-minute sampling. The results are recorded in Table 7. As for 30-minute sampling, we find that the performance of the AR improves relative to the HAR when the EDFs are used instead of the normal distribution assumption, and there are benefits to using the HAR relative to the AR when the normal distribution is used. In contrast with the results for AU in Table 5a, there is no evidence that the direct use of intraday returns (i.e., the MIDAS model) improves VaR forecasts in comparison to an AR model.

### 5.6. Explaining country-specific differences in forecast performance

Our results indicate the largest differences across models and methods are for CA. Graphs of the daily returns in the out-of-sample period (July 4, 2002 to October 27, 2003) along with 5% VaR forecasts from the AR and HAR models (computed assuming a normal distribution) shed some light on these findings. Fig. 3 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. Fig. 4 shows the 1-step ahead forecasts of realized volatility (and outturns) for all five countries. From this figure, it is clear that the good performance of the HAR model VaR forecasts for CA 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 relative to either the AR or MIDAS models. It is also apparent from Fig. 4 that currencies other than CA do not show such a clear level shift, or such a clear distinction between the volatility forecasts of the models.

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 reference to the in-sample period estimates of the AR and HAR, which are given in the Appendix A. The parameter estimates recorded there correspond directly to the model parameters defined by Eqs. (3) and (5) for the AR and HAR, respectively. Also reported are the estimates for the MIDAS model given in Eq. (4)). Consider a fixed-forecasting scheme, and suppose that there has been a level shift in the unconditional mean of  $\log(RV_t^{1/2})$  from its in-sample value of  $v$  to  $v'$ . The average values of the forecasting functions of the AR and HAR models sufficiently long after the shift has occurred can be approximated by substituting RHS terms in  $\log(RV_t^{1/2})$  by  $v'$ . Coupled with the in-sample parameter estimates, the average values of the forecast functions are (for Canada):

$$\begin{aligned} \text{AR} : & \quad -2.70 + 0.53v' \\ \text{HAR} : & \quad -1.62 + 0.72v'. \end{aligned}$$

The greater weight on  $v'$  in the HAR model means that its volatility forecasts will track the shift in the underlying level of volatility better than the forecasts of the AR, because it adjusts faster to the level shift.

Although not shown in Fig. 3, the use of the EDF (based on rolling windows) to compute VaR instead of the assumption of normality improves the performance of both models, but especially that of the AR model (in keeping with the results in Table 4a).

In general, we find that a simple model for realized volatility (an AR) 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 the Canadian dollar, 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 exchange rate volatility at daily, weekly and fortnightly horizons. We find that the method of parameterizing intraday returns implicit in the step-function MIDAS (i.e., the 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. 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 in general.

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 autoregression) combined with the assumption of a normal distribution for expected future returns yields reasonable VaR estimates for the majority of the currencies in our sample. The exception is the Canadian dollar, and we explain the different findings for this currency in terms of a specific structural break in the underlying level of volatility in the out-of-sample period.

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, or even an assumption of normality, will in 'normal times' generate competitive Value-at-Risk forecasts with reasonable coverage rates, although when there are structural shifts models such as the HAR may fare better.

**Appendix A. Estimates of AR, MIDAS and HAR for each currency and horizon**

In-sample period: January 4, 1999 to July 3, 2002 (common sample: same sample size for each h and model).

AR			MIDAS			HAR					
<i>Australian dollar</i>											
	<i>h=1</i>	<i>h=5</i>	<i>h=10</i>		<i>h=1</i>	<i>h=5</i>	<i>h=10</i>		<i>h=1</i>	<i>h=5</i>	<i>h=10</i>
$\psi_0$	−1.19	−.675	−.576	$\beta_0$	.004	.409	.509	$\beta_0$	−.635	.064	.218
$\psi_1$	.276	.204	.169	$\beta_1$	.735	.664	.627	$\beta_D$	.159	.091	.056
$\psi_2$	.120	.151	.133	$\theta_1$	−.017	.0101	.002	$\beta_W$	.340	.263	.218
$\psi_3$	.049	.111	.116	$\theta_2$	.0001	.000	.000	$\beta_M$	.328	.488	.527
$\psi_4$	.122	.120	.114								
$\psi_5$	.193	.107	.111								
$R^2$	.32	.44	.44	$R^2$	.31	.44	.44	$R^2$	.32	.48	.50
<i>Canadian dollar</i>											
	<i>h=1</i>	<i>h=5</i>	<i>h=10</i>		<i>h=1</i>	<i>h=5</i>	<i>h=10</i>		<i>h=1</i>	<i>h=5</i>	<i>h=10</i>
$\psi_0$	−2.70	−2.48	−2.59	$\beta_0$	−1.97	−1.74	−2.06	$\beta_0$	−1.62	−1.19	−1.32
$\psi_1$	.204	.133	.105	$\beta_1$	.495	.412	.153	$\beta_D$	.131	.063	.049
$\psi_2$	.162	.079	.073	$\theta_1$	.010	−.003	.010	$\beta_W$	.215	.139	.062
$\psi_3$	.011	.056	.048	$\theta_2$	.000	.000	.000	$\beta_M$	.371	.439	.447
$\psi_4$	.035	.070	.052								
$\psi_5$	.115	.077	.056								
$R^2$	.12	.17	.16	$R^2$	.11	.17	.15	$R^2$	.12	.22	.24
<i>Euro</i>											
	<i>h=1</i>	<i>h=5</i>	<i>h=10</i>		<i>h=1</i>	<i>h=5</i>	<i>h=10</i>		<i>h=1</i>	<i>h=5</i>	<i>h=10</i>
$\psi_0$	−1.54	−1.20	−1.12	$\beta_0$	−.369	−.225	−.301	$\beta_0$	−.949	−.285	−.009
$\psi_1$	.185	.184	.146	$\beta_1$	.678	.576	.514	$\beta_D$	.057	.089	.055
$\psi_2$	.133	.138	.116	$\theta_1$	−.005	.010	−.0004	$\beta_W$	.451	.217	.132
$\psi_3$	.124	.108	.102	$\theta_2$	.000	.000	.000	$\beta_M$	.306	.468	.572
$\psi_4$	.059	.076	.087								
$\psi_5$	.197	.089	.091								
$R^2$	.23	.34	.33	$R^2$	.22	.33	.32	$R^2$	.23	.39	.43
<i>British Pound</i>											
	<i>h=1</i>	<i>H=5</i>	<i>h=10</i>		<i>h=1</i>	<i>h=5</i>	<i>h=10</i>		<i>h=1</i>	<i>h=5</i>	<i>h=10</i>
$\psi_0$	−1.83	−1.32	−1.29	$\beta_0$	−.414	−.069	−.246	$\beta_0$	−1.12	−.599	−.463
$\psi_1$	.239	.165	.147	$\beta_1$	.686	.617	.544	$\beta_D$	.133	.061	.051
$\psi_2$	.099	.115	.103	$\theta_1$	−.011	−.002	.001	$\beta_W$	.322	.334	.256
$\psi_3$	.090	.106	.099	$\theta_2$	.0001	.000	.000	$\beta_M$	.338	.336	.384
$\psi_4$	.044	.097	.090								
$\psi_5$	.189	.113	.097								
$R^2$	.20	.36	.36	$R^2$	.20	.37	.36	$R^2$	.21	.39	.40
<i>Japanese Yen</i>											
	<i>h=1</i>	<i>H=5</i>	<i>h=10</i>		<i>h=1</i>	<i>h=5</i>	<i>h=10</i>		<i>h=1</i>	<i>h=5</i>	<i>h=10</i>
$\psi_0$	−1.65	−1.36	−1.39	$\beta_0$	−.443	−.297	−.460	$\beta_0$	−1.06	−.548	−.399
$\psi_1$	.266	.200	.157	$\beta_1$	.671	.569	.495	$\beta_D$	.169	.119	.080
$\psi_2$	.128	.120	.099	$\theta_1$	−.017	−0.010	−.009	$\beta_W$	.335	.207	.096
$\psi_3$	.094	.097	.084	$\theta_2$	.0001	.0001	.0001	$\beta_M$	.292	.401	.511
$\psi_4$	.051	.073	.073								
$\psi_5$	.142	.070	.080								
$R^2$	.24	.32	.29	$R^2$	.24	.33	.30	$R^2$	.25	.36	.37

## References

- Andersen, T.G., 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., Diebold, F.X., 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., Labys, P., 2003a. Great realizations. *Risk* 13, 105–108.
- Andersen, T.G., Bollerslev, T., Diebold, F.X., Labys, P., 2003b. Modeling and forecasting realized volatility. *Econometrica* 71, 579–625.
- Andersen, T.G., Bollerslev, T., 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., Meddahi, N., 2006. Realized volatility forecasting and market microstructure noise, Department of Finance, Kellogg School of Management, Northwestern University, mimeo.
- Bandi, F., Russell, J., 2005. Separating microstructure noise from volatility. *Journal of Financial Economics* 79, 655–692.
- Bandi, F., Russell, J., in press. Microstructure noise, realized variance, and optimal sampling. *Review of Economic Studies*.
- Bao, Y., Lee, T.H., Saltoğlu, B., 2004. A test for density forecast comparison with applications to risk management, UC Riverside, mimeo.
- Barndorff-Nielsen, O.E., 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., 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., Shephard, N., 2003. Realised power variation and stochastic volatility. *Bernoulli* 9, 243–265.
- Barndorff-Nielsen, O.E., 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., Taylor, S.J., 2001. Forecasting S and 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., Nelson, D.B., 1994. ARCH models. In: Engle, R.F., McFadden, D. (Eds.), *The Handbook of Econometrics*, vol. 4. North-Holland, Amsterdam, pp. 2959–3038.
- Christoffersen, P.F., 1998. Evaluating interval forecasts. *International Economic Review* 39, 841–862.
- Clements, M.P., Hendry, D.F., 2006. Forecasting with breaks. In: Elliott, G., Granger, C.W.J., Timmermann, A. (Eds.), *Handbook of Economic Forecasting*, Vol.1. *Handbook of Economics*, vol. 24. Elsevier, North-Holland, pp. 605–657.
- Comte, F., 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, mimeo.
- Corsi, F., Kretschmer, U., Mitnik, S., Pigorsch, C., 2005. The volatility of realized volatility. Discussion paper, no. 2005/33, Centre for Financial Studies, University of Frankfurt.
- Diebold, F.X., Mariano, R.S., 1995. Comparing predictive accuracy. *Journal of Business and Economic Statistics* 13, 253–263.
- Engle, R.F., 1982. Autoregressive conditional heteroscedasticity, with estimates of the variance of united kingdom inflation. *Econometrica* 50, 987–1007.
- Forsberg, L., Ghysels, E., 2004. Why do absolute returns predict volatility so well? University of North Carolina, mimeo ([www.unc.edu/eghysels](http://www.unc.edu/eghysels)).
- Ghysels, E., Santa-Clara, P., Valkanov, R., 2004. The MIDAS touch: Mixed data sampling regression models. Chapel Hill, mimeo.
- Ghysels, E., Santa-Clara, P., 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., Valkanov, R., 2006. MIDAS regressions: Further results and new directions. *Econometric Reviews* 26, 53–90.
- Giacomini, R., Komunjer, I., 2005. Evaluation and combination of conditional quantile forecasts. *Journal of Business and Economic Statistics* 23, 416–431.
- Giot, P., 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., Kamstra, M., 1989. Interval forecasting: an analysis based upon ARCH-quantile estimators. *Journal of Econometrics* 40, 87–96.
- Hansen, P.R., Lunde, A., 2006. Realized variance and microstructure noise. *Journal of Business and Economic Statistics* 24, 127–161.
- Koopman, S.J., Jungbacker, B., 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., Granger, C.W.J., 2003. Forecasting volatility in financial markets: A review. *Journal of Economic Literature* 41, 478–539.
- Riskmetrics, 1995. Technical Document, JP Morgan. New York, USA.
- West, K.D., 2006. Forecasting evaluation. In: Elliott, G., Granger, C.W.J., Timmermann, A. (Eds.), *Handbook of Economic Forecasting*, vol.1. *Handbook of Economics*, vol. 24. Elsevier, North-Holland, pp. 99–134.
- White, H., 2000. A reality check for data snooping. *Econometrica* 68, 1097–1126.
- Zhang, L., Mykland, P.A., Ait-Sahalia, Y., 2005. A tale of two time scales: Determined integrated volatility with noisy high-frequency data. *Journal of the American Statistical Association* 100, 1394–1411.