AUCO Czech Economic Review 4 (2010) 295–314
Acta Universitatis Carolinae Oeconomica

Forecasting the Quantiles of Daily Equity Returns Using Realized Volatility: Evidence from the Czech Stock Market

Vít Bubák*

Received 1 March 2010; Accepted 3 May 2010

Abstract In this study, we evaluate the quantile forecasts of the daily equity returns on three of the most liquid stocks traded on the Prague Stock Exchange. We follow the recent findings that consider the potential value of intraday information for volatility forecasting and, instead of proxying volatility using daily squared returns, we use both the intraday returns as well as their lower frequency aggregate (realized volatility) to forecast volatility and ultimately the quantiles of the distributions of future returns under different scenarios. We find that a simple autoregressive model for realized volatility together with the assumption of a normal distribution for expected returns results in VaR forecasts that are no worse than those based on other models (HAR, MIDAS) and/or other methods of computing the distribution of future returns. In fact, similar results obtain across the different forecast horizons and at both 2.5% and 5% VaR levels despite superior performance of HAR model in out-of-sample volatility forecasts.

Keywords Intraday data, heterogeneous autoregressive model, mixed data sampling model, realized volatility, Value-at-Risk
JEL classification C01, C32, C53, G32

1. Introduction

The research into reliable and accurate risk-management methodologies has for the last twenty years represented an active and growing area of financial econometrics. Inspired in the financial markets where every day, both financial and non-financial institutions face the difficult task of estimating the extent of market risk exposure so as to make better informed and more efficient capital decisions in the future, the importance of risk-management research has only intensified with the recent turbulent economic events.

The central focus of this study is the most common measure of risk in use today, the Value-at-Risk (VaR). Developed in the 1990s to quantify and assess the market risk exposure, the widespread popularity of the VaR owes mainly to its adoption as the First Pillar in the Basle II agreements (Basel Committee 2005) that regulates the total

* Charles University, Faculty of Social Sciences, Institute of Economic Studies, Opletalova 26, CZ-110 00 Prague, Czech Republic; Université Paris I, Centre d'Economie de la Sorbonne, Panthéon-Sorbonne, 106-112 bld de l'Hôpital 433, 756 47 Paris, France. E-mail: vit.bubak@malix.univ-paris1.fr, Phone: +420 222 112 330.

minimum capital requirements for credit, market and operational risk faced by a bank.[1]

We adopt a standard approach to measuring and estimating the VaR, structured as one involving both the choice of a model to forecast volatility and of the method of obtaining the quantile estimates from the volatility forecasts; unlike many of the studies, however, we depend solely on the use of high-frequency returns in order to forecast volatility. In other words, instead of proxying volatility using (daily) squared returns, we exploit the information inherent in intraday returns to construct the volatility forecasts and ultimately the quantiles of the distributions of future returns.

Our motivation stems from the recent financial econometrics literature that has shown a potential value of using both the lower-frequency aggregates of intraday returns such as realized volatility or realized power variation (e.g. Andersen et al. 2003; Barndorff-Nielsen and Shephard 2003), or even intraday returns themselves (Ghysels et al. 2006; Ghysels et al. 2007) for volatility forecasting. With this literature in mind and, understanding that there are a number of other models that we could potentially use, we choose to employ the Heterogeneous Autoregressive model—or HAR—of Corsi (2009) that makes use of the intraday returns in aggregated form, and the MIxed Data Sampling model—or MIDAS—of Ghysels et al. (2006) that is known for its ability to project the dependent variable directly onto intraday returns. Both of these models have been found to predict volatility reasonably well, especially at lower frequencies. In addition, we also consider a simple autoregressive model as a benchmark.

Following a brief evaluation of the relative forecasting performance of the alternative volatility models, we use the predicted volatility to construct the quantile forecasts. As in Clements et al. (2008), we consider two ways of calculating the quantile forecasts. In particular, we assume either a specific distribution for the predictive density of future returns (Normal and Student-*t*) or obtain the predicted quantiles directly, using the empirical distribution function of the returns standardized with volatility forecasts. We examine the performance of the two methods relative to the benchmark Normal and AR cases.

Our analysis is based on intraday returns on three of the most liquid stocks traded on the Prague Stock Exchange, a post-emerging equity market in the Central and Eastern European (CEE) region. To the best of our knowledge, this study is the first in the CEE region to use high-frequency intraday data in order to estimate the quantiles of (daily) equity returns and second only to Žikeš (2008) to forecast realized volatility. In this respect, we contribute to a growing literature that investigates the forecasting power of empirical models designed to make use of intraday data and their usefulness in practice.

Our results show that a simple autoregressive model for realized volatility together with the assumption that the expected future returns follow a normal distribution leads to the quantile forecasts that are at least as good as those obtained from the other models. In fact, the same performance of the AR model holds true across the equity returns analyzed, regardless of the forecast horizon. Furthermore, the results from the quantile forecasts obtain despite the fact that HAR models seem to perform better than

[1] See Jorion (2005) for an extensive review of major applications of the VaR approach.

either AR or MIDAS models in the out-of-sample volatility forecasting exercise.

The remainder of the paper is organized as follows. In the next section, we briefly describe the data and discuss how we construct the returns. Following a note on the choice of the appropriate sampling frequency, we also discuss the derivation of the volatility measures employed in the empirical part of the study. The section concludes with the descriptive statistics. In Section 3, we describe the volatility forecasting models and, in Section 4, the methodology behind the construction and evaluation of the volatility and quantile forecasts. Finally, in Section 5 we present the empirical results. Section 6 concludes.

2. Intraday returns and volatility

2.1 Intraday returns

The empirical part of our study is based on the intraday price data for three of the most liquid companies traded on SPAD segment of the Prague Stock Exchange:[2] ČEZ (further denoted as CEZ), the largest electricity producer in the country, Telefónica O2 Č.R. (TEF), a telecommunications company, and Komerční Banka (KOB), the third largest bank in the country.[3] The sample runs for the period from January 3, 2000 to December 30, 2008 yielding a total of $N = 2,246$ trading days. A shorter dataset is examined in Žikeš (2008) and we refer the reader to the latter study for any details that she might find missing here.

Prior to the construction of the intraday returns, we perform a basic cleaning procedure following Barndorff-Nielsen et al. (2008). As part of the procedure, we also remove all observations outside of regular trading hours which, until April 1, 2008, ran from 9:30am till 4:00pm Central European Time. Exclusion of the overnight information remains a standard practice in the volatility modeling literature when only intraday information needs to be considered (see e.g. Bollerslev et al. 2009). In particular, it allows to avoid distortions associated with lower liquidity during the non-business trading hours. This said, towards the end of the study we also consider the robustness of the results when the overnight information is included.

We define the intraday (raw) returns as the first difference of the logarithms of the mid-points of the best bid (B) and the best ask (A) prices, sampled along a 30 minute time grid. Formally, at any given time $j = 1,\ldots,m$ of day t, where m denotes the number of intraday returns, an intraday (raw) return is defined as $r_{t,j} = 100(p_{t-1+j/m} - p_{t-1+(j-1)/m})$, where $p = \ln P$ and $P = (B+A)/2$. Given the 30 minute sampling frequency and six-and-a-half regular trading hours, we have $m = 13$ intraday returns. This gives us a total of $21,198$ intraday returns for the whole sample.

The choice of the 30 minute sampling frequency is, of course, not arbitrary. It is

[2] SPAD, or Stock and Bond Market Support System, is a price-driven trading system run by PSE based on the activity of market makers. A detailed description of the SPAD is provided by Hanousek and Podpiera (2004). Additional information can be found in Bubák and Žikeš (2006).

[3] As of March 1, 2009, the combined capitalization of the three stocks was nearly 65% of the PX index, the main market index of the PSE. Note also that prior to May of 2005, Telef ónica O2 Č.R. (Bloomberg: `SPTT CP`) traded as Český Telecom.

well known that the high-frequency returns are subject to a host of market microstructure features including bid-ask bounce effects, discrete price observations, infrequent and nonsynchronous trading, et cetera, which may all distort the properties of returns. Of immediate consequence of the presence of noise for any study that uses high-frequency returns to estimate the ex-post daily volatility, including our own, is that a simple estimator of daily volatility such as realized volatility (discussed further in the text) may no longer be consistent. Bandi and Russell (2005) provide a simple technique to identifying both the genuine time-varying volatility of the unobservable returns and the variance induced by the microstructure noise.

We follow the findings of Žikeš (2008) who discusses the choice of the sampling frequency in detail. In particular, Žikeš (2008) suggests using a 30 minute sampling frequency as a reasonable compromise between the noise-induced bias introduced to the realized volatility estimator by sampling more frequently and the loss of information from sampling at frequencies lower than 30 minutes. The same estimator is also shown to behave better than a microstructure-robust estimator of Zhou (1996) which, even at the 30 minute frequency, remains biased relative to the former and only experiences an increase in its variance at the frequencies beyond 30 minutes.

2.2 Volatility measures

Andersen and Bollerslev (1998) were perhaps the first to point out that high-frequency data can be used to form both more accurate and meaningful *ex-post* volatility measurements. Andersen et al. (2001) elaborated on these findings when they formally showed that the ex-post volatility may be estimated to any degree of accuracy simply by summing sufficiently high-frequency returns within a day. The corresponding measure, termed (daily) realized volatility (Andersen et al. 2000), is then defined as a cumulative sum of squared intraday returns,

$$RV_{t,t+1} = \sum_{j=1}^{m} (r_{t,j})^2, \tag{1}$$

where $r_{t,j}$ represents an (intraday) return obtained for time j of day t.

In the absence of microstructure noise, (1) can be shown to be consistent for the so called integrated variance, a natural albeit unobservable measure of volatility, plus a possible jump component (Andersen et al. 2001). As already noted, in the presence of an equally unobservable microstructure noise the realized volatility becomes biased, with the variance growing larger the higher the sampling frequency. Barndorff-Nielsen and Shephard (2002) study the asymptotic properties of realized volatility and present conditions under which it is also an unbiased estimate. Andersen et al. (2003) discuss the theoretical framework underlying the construction and general properties of realized volatility.

In our study, we use both a daily version of RV, defined in (1), as well as its five-day (weekly) and ten-day (bi-weekly) versions obtained by summing (1) over five-day and ten-day periods ($D = 5, 10$), respectively, as follows

$$RV_{t,t+D} = \sum_{j=1}^{D} RV_{t+j-1,t+j}. \tag{2}$$

Several other classes of empirical processes based on high-frequency data and related to volatility have been found to have a predictive power. We choose to examine the predictive performance of the realized absolute variation, *RAV*, in addition to *RV* itself. First introduced by Barndorff-Nielsen and Shephard (2004), the *RAV* is constructed in terms of absolute values of intraday returns as

$$RAV_{t,t+1} = \mu^{-1} m^{-1/2} \sum_{j=1}^{m} |r_{t,j}|, \tag{3}$$

where $\mu \equiv \sqrt{2/\pi}$ denotes the mean of the absolute value of standard Normal random variable, $E(|Z|)$. The fact that an absolute value (and not a square) function is used in the construction of *RAV* provides a clue as to why the *RAV* has a potential to improve the volatility modeling.

First, it has been long recognized that the absolute value returns reflect stronger persistence than squared returns, therefore providing a potentially better signal for volatility (see e.g. Ding et al. 1993).[4] Second, the absolute returns are relatively less sensitive to large price movements and as such may provide more accurate predictions during the periods with jumps. In fact, Barndorff-Nielsen and Shephard (2004) show that under certain mild conditions, the *RAV* is not affected by the jump component of the underlying measure of intraday variation. To some extent, such *immunity* is relevant even to our study, as Žikeš (2008) provides evidence of the presence of rare jumps in the same, albeit shorter set of data. Finally, Ghysels et al. (2006) and Forsberg and Ghysels (2007) show that regressions involving the absolute values of high-frequency returns as explanatory variables can improve forecasts of realized volatility at lower frequencies.

2.3 Descriptive statistics

We employ a logarithm of the square root of realized volatility as a dependent variable in all regressions. The logarithmic transformation has been commonly applied in many studies on realized volatility forecasting, including e.g. Forsberg and Ghysels (2007), Andersen et al. (2007) and Clements et al. (2008), and follows naturally from an observation that the distribution of the logarithmic realized volatility (or, logarithmic standard deviation, $\log RV^{1/2}$), can be well approximated by the Gaussian distribution. First noted by Andersen et al. (2001) and Andersen et al. (2003), the log-normality of realized volatility also suggests and greatly facilitates the use of standard linear approaches to modeling and forecasting the logarithmic realized volatilities.

[4] Liu and Maheu (2005) also mention that if higher order moments of returns (e.g. fourth moment) do not exist, the absolute value of returns will be more reliable as its variance is more likely to exist.

Table 1. Descriptive statistics

		Mean	S.D.	Skew.	Kurt.	Min	Max	JB	LB(20)
				A: Logarithmic realized SD					
CEZ	l_rv_1	0.195	0.585	0.024	3.308	−1.780	2.290	9.082	2,719
	l_rv_5	1.144	0.425	0.527	3.630	−0.006	2.724	141.1	11,258
	l_rv_{10}	1.527	0.385	0.665	3.704	0.631	2.934	211.7	17,706
KOB	l_rv_1	0.251	0.586	0.011	3.333	−1.648	2.333	10.44	3,748
	l_rv_5	1.193	0.433	0.515	3.606	−0.153	2.679	133.6	14,465
	l_rv_{10}	1.572	0.400	0.621	3.621	0.610	2.961	180.6	20,623
TEF	l_rv_1	0.099	0.788	−0.741	4.771	−4.169	2.243	499.1	13,628
	l_rv_5	1.054	0.667	−0.962	5.506	−2.216	2.530	934.2	24,730
	l_rv_{10}	1.441	0.626	−0.990	5.481	−1.833	2.829	942.4	29,226
				B: Standardized returns					
CEZ	s_r_1	0.102	1.070	−0.056	2.073	−2.573	2.617	81.62	29.24
	s_r_5	0.172	1.090	0.020	2.480	−2.998	3.101	25.44	2,569
	s_r_{10}	0.224	1.098	0.059	2.520	−2.799	3.227	22.89	7,024
KOB	s_r_1	0.042	1.047	−0.051	2.172	−2.676	3.007	65.19	26.72
	s_r_5	0.088	1.099	0.015	2.511	−3.053	3.514	22.47	2,778
	s_r_{10}	0.114	1.109	0.017	2.595	−2.993	3.489	15.47	6,384
TEF	s_r_1	−0.006	1.033	0.080	2.163	−2.588	2.693	67.89	16.55
	s_r_5	−0.001	1.077	0.026	2.416	−2.732	3.012	32.21	2,818
	s_r_{10}	−0.006	1.099	0.034	2.544	−3.248	3.064	19.90	6,519

Note: JB contains Jarque-Bera test statistic for the null hypothesis of normality of the series; i.e., Skew. = 0 and Kurt. = 3 (5% crit. value is 5.99). The column labeled LB(20) contains Ljung-Box test statistics for the null hypothesis of no autocorrelation up to lag 20 in the series (5% crit. value is 31.41). The sample runs from January 3, 2000 to December 30, 2008 (2,246 observations).

Panel A of Table 1 reports basic descriptive statistics for the logarithms of the square root of realized volatility, $\log(RV_{t,t+H}^{1/2})$, denoted in the table as l_rv_H for the particular horizons of $H = (1,5,10)$ days. The statistics make it evident that most of the returns series are not normally distributed; indeed, of the nine series in this group, only two series (CEZ, KOB) seem to be reasonably close to being normal in case of $H = 1$, owing largely to the relatively small skewness and excess kurtosis of only around 0.3. Figure 1 (upper row) confirms this observation by plotting the kernel density estimates for the case of l_rv_1.

The Ljung-Box statistics indicate a presence of strong serial correlation in the realized volatilities for all three stocks, reflecting the long-memory property of the realized volatility. Žikeš (2008) discusses the basic distributional and dynamic characteristics of the three stock returns and realized volatilities in greater detail and provides also a graphical confirmation of an apparent hyperbolic decay characteristic of a long-range serial dependence present in the series.

Panel B presents descriptive statistics for one-day, five-day, and ten-day returns standardized with the appropriate square root of realized volatility, $r_{t,t+H}/RV_{t,t+H}^{1/2}$. Again, the series are denoted as s_r_H for simplicity. First note that, although much

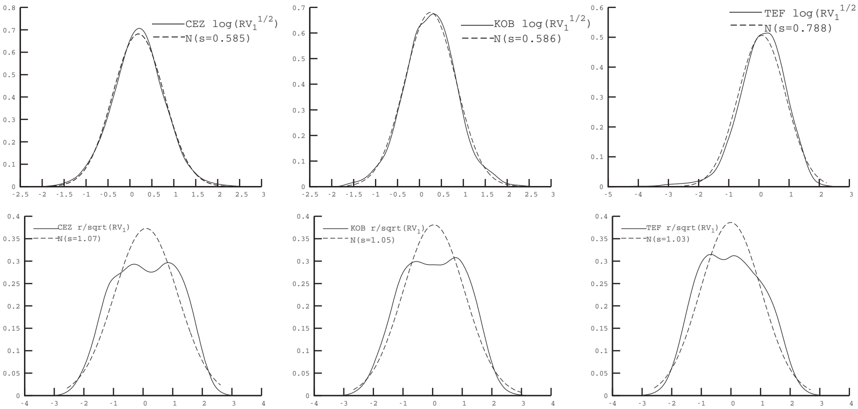

Figure 1. Kernel density estimates for l_rv_1 and s_r_1

more symmetric than l_rv_H, the null hypothesis of normality of standardized returns is similarly rejected at any reasonable level of significance for most of the series.

The kernel density estimates of s_r_1, shown in Figure 1 (lower row), illustrate additional points related to the normality of the standardized returns. For example, it is the relatively low sampling frequency ($m = 13$) that drives the fourth moment of the distribution of standardized returns well below its asymptotic (normal) value of 3 as much as it restricts the support of the distribution to the [-$\sqrt{m}$,$\sqrt{m}$] bounds (about $[-3.6, 3.6]$ in our study). Perceptibly thin tails evident in the particular shapes of the kernel density estimates are also given by construction; see Andersen et al. (2007).[5]

3. Volatility forecasting models

In this section we discuss the three regression models used in our study to predict the future realized volatility: the heterogeneous autoregressive model (HAR) of Corsi (2009), the MIxed Data Sampling (MIDAS) model of Ghysels et al. (2004), and a plain-vanilla AR model that we use as a benchmark. We specify the AR model first, followed by HAR and MIDAS models.

The AR specification

The simplest forecasting model model employed in this study is a simple autoregression of the logarithm of the square root of realized volatility on its past values:

$$\log RV_{t,t+H}^{1/2} = \alpha_0 + \sum_{d=1}^{p} \alpha_d \log RV_{t-d,t-d+1}^{1/2} + \varepsilon_{t+H}, \tag{4}$$

[5] Unlike the center of the distribution that would tend to standard normal as $m \rightarrow \infty$, however, the tails would remain thin and only *close* to normal under the same circumstances.

where p represents the number of lags. In particular, we let $p = 5$, so that we use the information of the previous five days to forecast the dependent variable.

The HAR specification

The heterogeneous autoregressive model of Corsi (2009) considers the previous one-day, five-day (one week), and twenty-day (1 month) realized volatility to predict the current value of the dependent variable. Formally, the model is given by the following equation:

$$\log RV_{t,t+H}^{1/2} = \beta_0 + \beta_1 \log \widetilde{RV}_{t-1,t}^{1/2} + \beta_2 \log \widetilde{RV}_{t-5,t}^{1/2} + \beta_3 \log \widetilde{RV}_{t-20,t}^{1/2} + \varepsilon_{t+H}, \quad (5)$$

where the regressors on the right hand side of (5) are defined as (the logarithms of) the normalized sums of the past one-period realized volatilities,

$$\widetilde{RV}_{t-h,t}^{1/2} = h^{-1} \sum_{k=1}^{h} RV_{t-k,t-k+1}^{1/2}, \quad (6)$$

where $h = (1, 5, 20)$ denotes the number of days over which the normalized sum is constructed. In the similar fashion, in the empirical part of the study we also employ realized absolute variation, (3), to estimate the dependent variable. The corresponding regressors on the right hand side of (5) are then constructed as in (6), i.e. as (normalized) sums of the past one-period realized absolute variations.

Originally developed with the purpose of modeling and forecasting the time series behavior of volatility, the HAR model has been shown to perform better or at least as well as other common volatility models, including ARFIMA and GARCH, be it in both in- and out-of-sample volatility exercises (see Žikeš 2008 for the case of Czech data). The structure of the HAR model, namely the short-term and the long-term realized volatility components that enter the model as regressors, also lends it an intuitive interpretation; specifically, the latter components allow it to account for the different reaction times of various market participants to the arrival of news, hence directly relating the corresponding long- and short-term volatility patterns over time. This way, the HAR model contrasts with the *true* long-memory models such as ARFIMA and FIGARCH which, other shortcomings aside, completely lack clear economic interpretation (Corsi 2009). For various recent applications of the HAR model, see e.g. Andersen et al. (2007), Forsberg and Ghysels (2007), or Chung et al. (2008).

The MIDAS specification

The particular version of the MIDAS regression model we use can be written as:

$$\log RV_{t,t+H}^{1/2} = \mu_H + \phi_H \log \left[\sum_{k=0}^{k^{\max}} b_H(k,\theta) L^{k/m} r_t^2 \right]^{1/2} + \varepsilon_{t+H}, \quad (7)$$

where ϕ_H is the scale parameter, $b_H(k,\theta)$ are the lag coefficients (weights) defined further below, and $L^{1/m}$ is a lag operator with the property that $L^{1/m} r_t^2 = r_{t-k/m}^2$. In other words, we are projecting the dependent variable onto past squared returns sampled at intraday frequency m.[6] In our study, we let $k^{\max} = 65$. Consequently, previous five days of intraday returns $(5 \cdot m = 65)$ enter the right hand side of (7), effectively providing the same information as in the case of (4) with $p = 5$.

We parametrize $b_H(k;\theta)$, $\theta = [\theta_1, \theta_2]$, as an Exponential Almon Lag,

$$b_H(k;\theta) = \frac{e^{\theta_1 k + \theta_2 k^2}}{\sum_{k=0}^{k^{\max}} e^{\theta_1 k + \theta_2 k^2}}. \tag{8}$$

The exponential parametrization has been known to provide a significant flexibility, as illustrated for a two-parameter case in Ghysels et al. (2005).[7] We leave further details of the model to the studies of Ghysels et al. (2004) and Ghysels et al. (2007).

As a relatively simple, parsimonious, yet flexible framework that allows for integration of high-frequency time-series data into forecasting lower frequency dependent variables, the MIDAS model has seen a wide range of applications, especially in the domain of financial and macroeconomic analysis. Although most of the empirical studies proved a satisfactory performance of the model based on the U.S. data (e.g. Forsberg and Ghysels 2007, Clements and Galvão 2008), the model has been shown to achieve good results also in the studies analyzing emerging markets returns (see e.g. Alper et al. 2008). Consequently, it is only natural to inquire about the ability of the model to perform well also in the Czech settings.

4. Quantile forecasts

The models discussed in the previous section are used to obtain the forecasts of the logarithm of the square root of realized volatility for each stock and forecast horizon. To recover the forecasts of realized standard deviation needed for the construction of the quantile forecasts (referred to interchangeably as VaR), we follow Forsberg and Ghysels (2007) and undo the logarithmic transformation simply as $[RV_{t,t+H}^{1/2}]_f = \exp([\log RV_{t,t+H}^{1/2}]_f)$, where f denotes a forecast.

At this point, it becomes instructive to recall the formal definition of Value-at-Risk (VaR). Let $r_{t,t+H}$ denote the sum of daily exchange rate returns from $(t+1)$ to $(t+H)$. Then the (conditional) VaR, $v_{t,t+H}(p)$, is implicitly defined as the level of return $r_{t,t+H}$ that is exceeded with probability α, $\alpha \in (0,1)$,

$$v_{t,t+H}(\alpha) \equiv \inf_{v}\{v : P_t(r_{t,t+H} \leq v | I_t) \geq 1 - \alpha\}, \tag{9}$$

where I_t is the information available at time t. From (9) it is clear that finding a v is equivalent to uncovering the conditional quantile of $r_{t,t+H}$,

$$v_{t,t+H}(\alpha) = F_{r_{t,t+H}}^{-1}(1 - \alpha | I_t), \tag{10}$$

[6] For reasons mentioned in Section (2.2) and similarly to Forsberg and Ghysels (2007), we also employ absolute returns, $|r_t|$, as an explanatory variable in (7).

[7] Indeed, if more than two parameters are used, the flexibility of the exponential specification surpasses in theory that of another specification widely employed in the literature, “Beta Lag” (Ghysels et al. 2007), that is based on a Beta function and only two parameters.

where $F^{-1}_{r_{t,t+H}}(\cdot|I_t)$ is an inverse of the conditional distribution function of $r_{t,t+H}$.

Assuming that the underlying exchange rate returns are unpredictable in the mean, i.e. $r_{t,t+H} = [RV^{1/2}_{t,t+H}]_f z_{t,t+H}$, where $z_{t,t+H}$ is *i.i.d*, we can then use the forecasts of realized standard deviation to obtain the H-step ahead quantile forecasts at level α as

$$\hat{v}_{t,t+H}(\alpha) = [RV^{1/2}_{t,t+H}]_f F_t^{-1}(1-\alpha). \tag{11}$$

The exposition presented so far raises a question of what distribution to use for F_t. Although not confirmed by the descriptive statistics (see Section 2.3), we consider a normal distribution as the first method of approximating the predictive distribution of standardized returns. In short, we let $F_t^{-1} = \Phi^{-1}$, where Φ is the *c.d.f.* of the (standard) normal distribution.

As a second method, we consider Student's t-distribution. Although not validated by our data, our choice is guided by Clements et al. (2008) who use it as an alternative to the normal distribution in order to approximate predictive distributions for the data that display descriptive statistics similar to our own. We use eight degrees of freedom to approximate the predictive distribution of standardized returns when employing the Student's t-distribution.[8]

In addition to using parametric distributions, we also employ an empirical distribution function (EDF) of the returns standardized by the forecasts of realized standard deviation, $r_{t,t+H}/[RV^{1/2}_{t,t+H}]_f$, to calculate the predicted quantiles; that is, we let

$$F_t^{-1} = Q^{-1}_{EDF}. \tag{12}$$

Clements et al. calculate the EDF using both recursive and rolling samples of previous forecasts and we adopt both formulations also in our study. In particular, we first estimate the model on the sample up to time T (in-sample period). We then use an integer part of rT, $r = 0.24$, observations of the standardized returns, constructed using the predicted values of realized standard deviation based on the period $(1-r)T+1,\ldots,T-H$ (window), to compute the first quantile forecast at $t = T$. The calculation of each additional quantile forecast from $t = T+1$ to $T+N-H$, where N is the size of the out-of-sample period and H is the forecast horizon, differs according to the forecasting scheme. While in the recursive scheme, the number of observations increases by one with each forecast (i.e., the window increases from $(1-r)T+1,\ldots,T-H$ to $(1-r)T+1,\ldots,T-H,\ldots,T-H+N$), in the rolling scheme, the number of observations used remains fixed at rT.

There are several ways to evaluate the accuracy of quantile forecasts. In our study, we employ what is known in the literature as a *tick* loss function (see e.g. Giacomini and Komujer 2005, Brownlees and Gallo 2008). If we let $e^i_{t,t+H} = r^i_{t,t+H} - v^i_{t,t+H}(\alpha)$, where i denotes the relevant model, we can write the tick loss function as

$$L^i_H(\alpha) = \left[\alpha - 1(e^i_{t,t+H} < 0)\right] \cdot e^i_{t,t+H}. \tag{13}$$

[8] Anticipating the results, we note that while reporting the results for eight degrees of freedom, we also performed the analysis with 4, 6, and 10 degrees of freedom; however the results did not differ substantially.

The forecast is said to be optimal if it minimizes $E_t\left[L^i_H(\alpha)\right]$. This expectation can then be estimated as $N^{-1}\sum_{t=1}^{n}\hat{L}^i_H(\alpha)$, where $\hat{L}^i_H(\alpha)$ is obtained using the forecasts $\hat{v}_{t,t+H}$.

It is evident that part of the definition of (13) is related to the measure of unconditional coverage, $\alpha - 1(e^i_{t,t+H} < 0)$, that is used in testing the null hypothesis that the number of violations implied by the model (actual coverage) equals to the nominal level α. Note, however, that in case of (13), the difference between the return and the predicted quantile, $e^i_{t,t+H}$, is weighted by α when $e^i_{t,t+H} > 0$ and by $(1-\alpha)$ otherwise.

Diebold and Mariano (1995) suggest a simple testing procedure to assess whether the differences between any two sets of forecasts are statistically significant. Given a set of forecasts (i, j), we first generate a loss differential

$$\delta_H(\alpha) = \left[\alpha - 1(e^i_{t,t+H} < 0)\right]\cdot e^i_{t,t+H} - \left[\alpha - 1(e^j_{t,t+H} < 0\right]\cdot e^j_{t,t+H}. \tag{14}$$

Under the null hypothesis of no difference in the predictive accuracy between the two sets of forecasts, or $H_0 : E(\delta_H(\alpha)) = 0$, the test statistic

$$S = \overline{\delta}_H(\alpha)\left(avar(\overline{\delta}_H(\alpha))\right)^{-1/2}, \tag{15}$$

where $\overline{\delta}_H(\alpha)$ is the average loss differential and $avar(\overline{\delta}_H(\alpha))$ is a consistent estimate of the asymptotic (long-run) variance of $\bar{\delta}$, is asymptotically distributed $N(0,1)$. We specify a one-sided test, implying that rejecting the null hypothesis renders the set of forecasts j more accurate than the set i.

5. Empirical Results

The methodological part of the study was organized so as to logically follow the two-stage procedure involved in the construction of quantile forecasts. Similarly, in this section we present the empirical results by first focusing on the relative performance of the volatility models in forecasting the realized volatility, followed by an evaluation of the various methods of obtaining the quantile forecasts.

We discuss the in-sample performance of the volatility models first, followed by a thorough description of the out-of-sample results. Table 2 presents the summary of an in-sample fit for the logarithm of the square root of realized volatility for the individual models, with HAR and MIDAS models further differentiated according to whether squared (HAR_{RV}, MIDAS_{RV}) or absolute returns (HAR_{RAV}, MIDAS_{RAV}) were used in the construction of the regressors.

Although not immediately clear, closer inspection of the results makes it evident that the HAR model based on RAV offers the best fit overall, followed by the MIDAS model estimated using absolute returns. Note, however, that if only squared returns were used in the estimation, the AR and MIDAS models would lead to better results, with an AR model providing nearly the same in-sample fit as computationally much more intensive MIDAS model, especially at longer forecast horizons.

Table 2. In-sample R^2

	$H=1$			$H=5$			$H=10$		
	CEZ	KOB	TEF	CEZ	KOB	TEF	CEZ	KOB	TEF
Levels									
AR(5)	0.222	0.236	0.599	0.232	0.307	0.651	0.233	0.278	0.615
HAR_{RV}	0.229	0.253	0.616	0.257	0.347	0.677	0.252	0.319	0.641
HAR_{RAV}	0.228	0.248	0.613	0.251	0.313	0.671	0.238	0.276	0.633
$MIDAS_{RV}$	0.227	0.253	0.607	0.246	0.346	0.653	0.243	0.320	0.613
$MIDAS_{RAV}$	0.226	0.250	0.607	0.249	0.305	0.651	0.240	0.265	0.612
Ratios									
HAR_{RV}	1.030	1.070	1.029	1.105	1.131	1.041	1.084	1.146	1.042
HAR_{RAV}	0.994	0.980	0.995	0.976	0.902	0.990	0.946	0.865	0.989
$MIDAS_{RV}$	0.996	1.019	0.990	0.979	1.106	0.974	1.048	1.160	0.968
$MIDAS_{RAV}$	0.995	0.989	0.999	1.013	0.877	0.996	0.986	0.827	0.999

Note: The lower part of the table presents the ratios of R^2 for the given model over the R^2 for the corresponding AR(5) model; i.e., $(R^2)^i_{H,s}/(R^2)^{AR}_{H,s}$, where i (s) denotes the model (stock), respectively.

Table 3. Out-of-sample RMSE comparison

	Fixed Scheme			Recursive Scheme			Rolling Scheme		
	CEZ	KOB	TEF	CEZ	KOB	TEF	CEZ	KOB	TEF
$H=1$									
AR(5)	0.444	0.459	0.476	0.440	0.456	0.476	0.437	0.455	0.476
HAR_{RV}	0.986	0.981	0.994	0.989	0.984	0.993	0.991	0.984	0.992
HAR_{RAV}	0.970	0.968	0.993	0.976	0.975	0.992	0.980	0.977	0.990
$MIDAS_{RV}$	1.051	1.031	1.078	1.052	1.034	1.071	1.050	1.027	1.070
$MIDAS_{RAV}$	1.017	1.021	1.073	1.032	1.026	1.068	1.034	1.021	1.067
$H=5$									
AR(5)	0.149	0.157	0.160	0.145	0.154	0.160	0.142	0.153	0.160
HAR_{RV}	0.978	0.949	0.966	0.984	0.956	0.967	0.993	0.955	0.965
HAR_{RAV}	0.977	0.934	0.986	0.988	0.944	0.983	0.994	0.945	0.976
$MIDAS_{RV}$	1.039	1.005	1.066	1.041	1.011	1.063	1.048	1.013	1.065
$MIDAS_{RAV}$	1.020	0.992	1.079	1.028	0.999	1.072	1.036	1.002	1.071
$H=10$									
AR(5)	0.105	0.113	0.110	0.103	0.110	0.110	0.102	0.110	0.110
HAR_{RV}	0.980	0.950	0.958	0.984	0.954	0.960	0.989	0.954	0.958
HAR_{RAV}	0.982	0.940	0.988	0.988	0.948	0.985	0.986	0.950	0.976
$MIDAS_{RV}$	1.142	1.006	1.016	1.028	1.012	1.013	1.038	1.017	1.012
$MIDAS_{RAV}$	1.014	1.002	1.044	1.018	1.007	1.036	1.023	1.013	1.030

Note: The RMSE is calculated as the square root of the sum of squared forecast errors divided by H. The initial (in-sample) estimation is based on 1,750 observations corresponding to the period from January 3, 2000, to January 25, 2007. There are 474 out-of-sample forecasts.

We use three forecasting schemes to obtain the out-of-sample (OOS) forecasts. To this end, we split the data into two periods: the in-sample period, $1,\ldots,T$, and the out-of-sample period, $T+1,T+2,\ldots,T+N$, where N is the number of OOS forecasts.

The in-sample period corresponds to the interval from January 3, 2000, to January 25, 2007, or 1,750 observations, while the out-of-sample period runs from January 26, 2007 to December 30, 2008. The division yields $N = 474$ OOS forecasts, or 21.3% of the entire sample.

In the fixed scheme, the in-sample period observations are used to estimate the parameters that remain fixed over the out-of-sample period. In the recursive scheme, an observation is added at each forecast origin so that, for example, the first four observations of the out-of-sample period, $T+1,\dots,T+4$, are used in addition to all observations of the in-sample period to obtain the forecast at $T+5$. Finally, the rolling scheme employs a fixed window of size T (length of the in-sample period), so that both the start and the end of the window successively increase by one observation prior to each OOS forecast.

The out-of-sample results—classified according to the forecasting scheme employed—are presented in Table 3. We use a root mean square error (RMSE) to compare the accuracy of the models. The use of a mean square error (MSE) has been a common practice in the recent volatility forecasting literature; see e.g., Forsberg and Ghysels (2007). As demonstrated by Patton (2006), the MSE loss function is also robust with regards to the volatility proxy used.

With the exception of the AR model, for each stock s, model i, and forecast horizon H, we report the ratio of RMSE over the RMSE for an AR model, $RMSE^{i}_{H,s}/RMSE^{AR}_{H,s}$, as a measure of the forecasting performance of the model relative to the AR benchmark. Thus, ratios below one indicate that a given model outperforms the benchmark.

The results show that HAR model is the best model overall, performing still better than the benchmark model at each forecast horizon. The use of absolute returns improves the performance of both HAR and MIDAS models marginally at $H = 1$, although the gains become generally smaller at longer forecast horizons. In case of the TEF stock, the models based on squared returns perform better yet than those based on absolute returns, especially at $H = 5$ and $H = 10$. For this reason, we will only work with the regressors based on squared returns when constructing the quantile forecasts.[9]

Additional conclusions emerge when we compare the performance of the models across the different forecasting schemes. Except for the AR model, there seems to be a small gain from updating the parameter estimates over the out-of-sample period. In fact, the rolling and recursive schemes perform marginally worse than the fixed scheme for both HAR and MIDAS models at each forecast horizon in case of two out of three stocks (CEZ and KOB). The TEF stock seems to benefit when the parameters remain fixed over the out-of-sample period.

The latter observations are confirmed by comparing the (expected) loss implied by the 5% quantile (VaR) forecasts when the standardized returns are assumed to follow normal distribution (Table 4). Similarly to the findings from the volatility forecasts, there seems to be a little advantage to using either a recursive or a rolling scheme in VaR forecasts, with only the TEF stock experiencing consistently albeit marginally smaller losses across the different models and horizons.

[9] Table A1 reports in-sample estimates from AR, HAR, and MIDAS models when the regressors are based on squared returns.

Table 4. Out-of-sample tick loss comparison: 5% VaR forecasts

	Fixed Scheme			Recursive Scheme			Rolling Scheme		
	CEZ	KOB	TEF	CEZ	KOB	TEF	CEZ	KOB	TEF
$H = 1$									
AR	12.39	13.76	9.646	12.50	13.89	9.601	12.58	13.87	9.499
HAR	12.43	13.89	9.287	12.56	14.04	9.277	12.68	14.10	9.221
MIDAS	12.64	13.93	9.656	12.71	14.04	9.591	12.80	13.99	9.454
$H = 5$									
AR	33.40	35.90	27.67	33.60	36.19	27.42	33.68	35.97	26.90
HAR	33.51	36.26	26.88	33.76	36.62	26.77	33.82	36.57	26.41
MIDAS	33.38	36.07	27.95	33.59	36.33	27.71	33.71	36.16	27.16
$H = 10$									
AR	47.65	50.28	40.42	47.80	50.69	39.97	47.60	50.19	39.06
HAR	47.65	50.69	39.32	47.91	51.23	39.07	47.76	50.95	38.36
MIDAS	47.45	50.52	40.52	47.64	50.85	40.09	47.54	50.48	39.11

Note: The entries are calculated as $100 \times (N^{-1} \sum_{t=1}^{N} \hat{L}_H^i(\alpha))$, where $\hat{L}_H^i(\alpha)$ denotes the tick loss function corresponding to model i, horizon H, and VaR level $\alpha = 5\%$. All models were estimated using realized volatility (or, squared intraday returns) as explanatory variables.

The latter observations are confirmed by comparing the (expected) loss implied by the 5% quantile (VaR) forecasts when the standardized returns are assumed to follow normal distribution (Table 4). Similarly to the findings from the volatility forecasts, there seems to be a little advantage to using either a recursive or a rolling scheme in VaR forecasts, with only the TEF stock experiencing consistently albeit marginally smaller losses across the different models and horizons.

Table 5. Evaluation of 5% VaR forecasts with normal distribution as a benchmark

	Fixed Scheme			Recursive Scheme			Rolling Scheme		
	CEZ	KOB	TEF	CEZ	KOB	TEF	CEZ	KOB	TEF
$H = 1$									
AR	1.116	1.112	1.088	1.358	1.342	1.383	1.365	1.382	1.355
HAR	1.116	1.112	1.087	1.383	1.382	1.444	1.393	1.431	1.432
MIDAS	1.116	1.114	1.098	1.271	1.256	1.236	1.315	1.288	1.256
$H = 5$									
AR	1.107	1.107	1.109	1.282	1.251	1.269	1.250	1.303	1.246
HAR	1.107	1.107	1.109	1.337	1.306	1.280	1.290	1.348	1.325
MIDAS	1.107	1.107	1.109	1.152	1.154	1.216	1.116	1.211	1.248
$H = 10$									
AR	1.109	1.109	1.110	1.109	1.184	1.244	1.133	1.235	1.347
HAR	1.109	1.109	1.109	1.138	1.218	1.320	1.168	1.284	1.428
MIDAS	1.109	1.109	1.110	1.081	1.049	1.193	1.047	1.266	1.114

We are now ready to evaluate the quantile forecasts. Table 5 reports the ratios of the expected tick loss when either Student's t or an empiricial distribution function (EDF) is used to approximate the predictive distribution of standardized returns, relative to the expected tick loss implied by the assumption of normal distribution. It turns out that neither the Student's t-distribution nor the EDF provide better VaR forecasts than the normal distribution at a standard 5% level. Qualitatively similar results (available on request) also hold for the 2.5% VaR forecasts. Note that we do not report the results of the Diebold-Mariano test here as none of the ratios is smaller than one.

To assess the accuracy of the quantile forecasts with respect to the volatility forecasting model employed, in Table 6 we present the ratios of the tick loss implied by the particular model (HAR, MIDAS) when either normal, Student's t or an empirical distribution is used over the tick loss when the corresponding AR model is employed. In other words, unlike in Table 5, we study the relative performance of the models under different distributional assumptions.

Table 6. Evaluation of 5% VaR forecasts with AR model as a benchmark

	Normal Dist.			Student-t (8) Dist.			Recursive EDF			Rolling EDF		
	CEZ	KOB	TEF	CEZ	KOB	TEF	CEZ	KOB	TEF	CEZ	KOB	TEF
$H = 1$												
HAR	1.003	1.009	0.963	1.003	1.010	0.962	1.021	1.039	1.005	1.024	1.045	1.017
	(0.422)	(0.417)	(0.422)	(0.422)	(0.417)	(0.421)	(0.432)	(0.418)	(0.434)	(0.416)	(0.396)	(0.438)
MIDAS	1.020	1.013	1.001	1.021	1.014	1.011	0.954	0.947	0.894	0.983	0.944	0.927
	(0.380)	(0.347)	(0.348)	(0.346)	(0.302)	(0.318)	(0.354)	(0.325)	(0.303)	(0.351)	(0.331)	(0.297)
$H = 5$												
HAR	1.003	1.010	0.971	1.003	1.010	0.971	1.046	1.054	0.979	1.035	1.045	1.033
	(0.401)	(0.394)	(0.401)	(0.401)	(0.394)	(0.401)	(0.444)	(0.435)	(0.417)	(0.395)	(0.390)	(0.372)
MIDAS	0.999	1.005	1.010	0.999	1.005	1.010	0.898	0.927	0.968	0.892	0.934	0.930
	(0.315)	(0.235)	(0.333)	(0.317)	(0.236)	(0.336)	(0.304)	(0.196)	(0.384)	(0.342)	(0.322)	(0.351)
$H = 10$												
HAR	1.000	1.008	0.973	1.000	1.008	0.972	1.026	1.037	1.032	1.031	1.048	1.031
	(0.408)	(0.400)	(0.401)	(0.409)	(0.400)	(0.402)	(0.350)	(0.353)	(0.465)	(0.361)	(0.351)	(0.371)
MIDAS	0.996	1.005	1.002	0.996	1.005	1.002	0.971	0.890	0.961	0.921	1.030	0.828
	(0.315)	(0.240)	(0.332)	(0.317)	(0.242)	(0.335)	(0.453)	(0.451)	(0.439)	(0.492)	(0.490)	(0.397)

Note: In the parentheses, we show the p-values of the Diebold-Mariano test for the null of no difference in the predictive accuracy between two sets of quantile forecasts.

The results indicate that the choice of the model is generally independent of the distribution assumed. In particular, AR model is the best model overall, irrespective of whether normal, Student's t or an empirical distribution is used to obtain the quantile forecasts. To be more precise, neither the HAR nor the MIDAS models yield the VaR forecasts which would be statistically more accurate than those coming from the AR model. It is also worth noticing that the use of either a recursive or a rolling EDF seems to improve the performance of the MIDAS model relative to the AR model. Nevertheless, in both cases, the differences in the predictive accuracy between the two sets of forecasts are again statistically insignificant at any reasonable level of significance.

A final note concerns the robustness of our results to the data that explicitly considers overnight information. Recall that this information was initially removed from the sample to avoid potential distortions associated with lower liquidity during the non-business hours, a practice that is otherwise common in the literature. Including the overnight returns increases the number of observations in the sample to 31,444.

The results from the analysis based on the dataset extended with overnight information (available on request) are quantitatively similar to the ones obtained previously. In particular, the expected tick loss calculations confirm a superior performance of the normal distribution relative to the other distributional assumptions. In addition and as in the earlier analysis, a simple autoregressive model is found to be the best model overall, regardless of the distribution employed to recover the quantile forecasts.

6. Conclusion

Inspired in recent volatility forecasting literature, our study makes use of the information inherent in the high-frequency intraday returns to forecast the quantiles of the distributions of future returns at different time horizons.

We follow a standard two-stage approach to the problem of construction of quantile forecasts: one in which we first obtain the forecasts of (the standard deviation of) realized volatility and next use these to generate the quantile forecasts via either a specific assumption on the distribution of expected future returns or using an empirical distribution of the expected future returns.

We employ the volatility forecasting models that make explicit use of the information intrinsic to intraday returns either by modeling the (non-parametric) realized volatility as a linear function of its own aggregate values (HAR model) or as a function of the higher-frequency intraday returns (MIDAS model). Accurate volatility forecasts obtained from these models are essential for the construction of the VaR forecasts as is the method of calculating the predictive distribution for the expected future returns.

Our findings relate to the intraday returns on three of the most liquid stocks traded on the Prague Stock Exchange. We show that a simple autoregressive model of realized volatility together with an assumption that the expected future returns follow a normal distribution leads to the quantile that are at least as accurate as those obtained from the other models employed. In particular, the superior performance holds true for all three stocks at either 5% or 2.5% VaR levels, regardless of the forecast horizon analyzed. Furthermore, the results from the quantile forecasts obtain despite the fact that HAR models seem to perform better than either AR or MIDAS models in the out-of-sample forecasting exercise. The findings from the analysis based on the data that incorporate overnight information confirm the former results.

Acknowledgment The author would like to thank two anonymous referees for helpful comments and suggestions. Financial support from the IES (Institutional Research Framework 2005-2010, MSM0021620841) is also gratefully acknowledged.

References

Alper, C. E., Fendoglu, S. and Burak, S. (2008). Forecasting Stock Market Volatilities Using MIDAS Regressions: An Application to the Emerging Markets. Munich Personal RePEc Archive No. 7460.

Andersen, T. G. and Bollerslev, T. (1998). Answering the Skeptics: Yes, Standard Volatility Models do Provide Accurate Forecasts. *International Economic Review*, 39(4), 885–905.

Andersen, T. G., Bollerslev, T. and Diebold, F. X. (2007). Roughing It Up: Including Jump Components in the Measurement, Modeling and Forecasting of Return Volatility. *Review of Economics and Statistics*, 89(4), 701–720.

Andersen, T. G., Bollerslev, T., Diebold, F. X. and Ebens, H. (2001). The Distribution of Stock Return Volatility. *Journal of Financial Economics*, 61(1), 43–76.

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(453), 42–55.

Andersen, T. G., Bollerslev, T., Diebold, F. X. and Labys, P. (2003). Modeling and Forecasting Realized Volatility. *Econometrica*, 71(2), 579–625.

Andersen, T. G., Bollerslev, T., and Dobrev, D. (2007). No-Arbitrage Semi-Martingale Restrictions for Continuous-Time Volatility Models subject to Leverage Effects, Jumps and i.i.d. Noise: Theory and Testable Distributional Implications. *Journal of Econometrics*, 138(1), 125–180.

Bandi, F. and Russell, J. R. (2005). Microstructure Noise, Realized Variance and Optimal Sampling. Graduate School of Business, University of Chicago, mimeo.

Barndorff-Nielsen, O. E., Hansen, P. R., Lunde, A. and Shephard, N. (2008). Realised Kernels in Practice: Trades and Quotes. *Econometrics Journal*, 4, 1–33.

Barndorff-Nielsen, O. E. and Shephard, N. (2002). Econometric Analysis of Realized Volatility and its Use in Estimating Stochastic Volatility Models. *Journal of the Royal Statistical Society B*, 64, 253–280.

Barndorff-Nielsen, O. E. and Shephard, N. (2003). Realized Power Variation and Stochastic Volatility Models. *Bernoulli*, 9(2), 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), 1–37.

Basel Committee (2001). International Convergence of Capital Measurement and Capital Standards: A Revised Framework. Basel, Basel Committee of the Bank for International Settlements.

Bollerslev, T., Kretschmer, U., Pigorsch, C. and Tauchen, G. (2009). A Discrete-Time Model for Daily S&P500 Returns and Realized Variations: Jumps and Leverage Effects. *Journal of Econometrics*, 150(2), 151–166.

Brownlees, C. T. and Gallo, G. M. (2008). On Variable Selection for Volatility Forecasting: The Role of Focused Selection Criteria. *Journal of Financial Econometrics*, 6(4), 513–539.

Bubák, V. and Žikeš, F. (2006). Trading Intensity and Intraday Volatility on the Prague Stock Exchange: Evidence from an Autoregressive Conditional Duration Model. *Czech Journal of Economics and Finance*, 56(5–6), 223–245.

Chung, H., Huang, C.-S. and Tseng, T.-C. (2008). Modeling and Forecasting of Realized Volatility Based on High-Frequency Data: Evidence from Taiwan. *International Research Journal of Finance and Economics*, 22, 178–191.

Clements, M. P. and Galvâo, A. B. (2008). Macroeconomic Forecasting With Mixed-Frequency Data: Forecasting Output Growth in the United States. *Journal of Business and Economic Statistics*, 26(4), 546–554.

Clements, M. P., Galvâo, A. B. and Kim, J. H. (2008). Quantile Forecasts of Daily Exchange Rate Returns from Forecasts of Realized Volatility. *Journal of Empirical Finance*, 15, 729–750.

Corsi, F. (2009). A Simple Approximate Long Memory Model of Realized Volatility. *Journal of Financial Econometrics*, 7(2), 174–196.

Diebold, F. X. and Mariano, R. S. (1995). Comparing Predictive Accuracy. *Journal of Business and Economic Statistics*, 13(3), 253–263.

Ding, Z., Granger, C. W. J. and Engle, R. F. (1993). A Long Memory Property of Stock Market Returns and a New Model. *Journal of Empirical Finance*, 1, 83–106.

Forsberg, L. and Ghysels, E. (2007). Why Do Absolute Returns Predict Volatility So Well? *Journal of Financial Econometrics*, 5, 31–67.

Ghysels, E., Santa-Clara, P. and Valkanov, R. (2004). The MIDAS Touch: Mixed Data Sampling Regression Models. Montreal, Center for Interuniversity Research and Analysis on Organizations, Working Paper No. 20.

Ghysels, E., Santa-Clara, P. and Valkanov, R. (2005). There is a Risk-Return Trade-Off after All. *Journal of Financial Economics*, 76, 509–548.

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(1–2), 59–95.

Ghysels, E., Sinko, A. and Valkanov, R. (2007). MIDAS Regressions: Further Results and New Directions. *Econometric Reviews*, 26(1), 53–90.

Giacomini, R. and Komujer, I. (2005). Evaluation and Combination of Conditional Quantile Forecasts. *Journal of Business and Economics Statistics*, 23(4), 416–431.

Hanousek, J. and Podpiera, R. (2004). Czech Experience with Market Maker Trading System. *Economic Systems*, 28(2), 177–191.

Jorion, P. (2006). *Value at Risk: The New Benchmark for Managing Financial Risk.* 3rd ed., New York, McGraw-Hill.

Liu, C. and Maheu, J. M. (2005). Modeling and Forecasting Realized Volatility: The Role of Power Variation. Toronto, University of Toronto, mimeo.

Patton, A. J. (2006). Volatility Forecast Comparison Using Imperfect Volatility Proxies. Sydney, University of Technology, Quantitative Finance Research Centre, Research Paper No. 175.

Žikeš, F. (2008). Properties of Czech Stock Returns and Volatility: Evidence from High-Frequency Data. London, Imperial College London, mimeo.

Zhou, B. (1996). High Frequency Data and Volatility in Foreign-Exchange Rates. *Journal of Business and Economic Statistics*, 14(1), 45–52.

Appendix

Table A1. In-sample regression estimates

		CEZ			KOB			TEF		
Model	Cf.	$H = 1$	$H = 5$	$H = 10$	$H = 1$	$H = 5$	$H = 10$	$H = 1$	$H = 5$	$H = 10$
AR	α_0	0.060^a	1.042^a	1.435^a	0.071^a	1.043^a	1.435^a	0.015	0.985^a	1.380^a
		(0.0127)	(0.0168)	(0.0165)	(0.0133)	(0.0173)	(0.0178)	(0.0138)	(0.0221)	(0.0249)
	α_1	0.343^a	0.181^a	0.143^a	0.279^a	0.185^a	0.156^a	0.401^a	0.283^a	0.226^a
		(0.0261)	(0.0162)	(0.0142)	(0.0276)	(0.020)	(0.0157)	(0.0297)	(0.0224)	(0.0222)
	α_2	0.105^a	0.088^a	0.081^a	0.141^a	0.126^a	0.104^a	0.134^a	0.151^a	0.130^a
		(0.0251)	(0.0130)	(0.0112)	(0.0271)	(0.014)	(0.0114)	(0.0268)	(0.0176)	(0.0179)
	α_3	0.011	0.055^a	0.062^a	0.082^a	0.086^a	0.075^a	0.100^a	0.128^a	0.118^a
		(0.0272)	(0.0132)	(0.0106)	(0.0267)	(0.014)	(0.0109)	(0.0254)	(0.0169)	(0.0172)
	α_4	0.074^b	0.078^a	0.071^a	0.080^a	0.070^a	0.067^a	0.135^a	0.120^a	0.120^a
		(0.0260)	(0.0134)	(0.0106)	(0.0248)	(0.014)	(0.0112)	(0.0276)	(0.0160)	(0.0166)
	α_5	0.102^a	0.080^a	0.073^a	0.092^a	0.075^a	0.063^a	0.115^a	0.109^a	0.130^a
		(0.0231)	(0.0158)	(0.0141)	(0.0258)	(0.017)	(0.0142)	(0.0235)	(0.0200)	(0.0202)
HAR	β_0	−0.004	0.970^a	1.374^a	−0.008	0.962^a	1.363^a	-0.071^a	0.892^a	1.291^a
		(0.0161)	(0.0207)	(0.0205)	(0.0179)	(0.0225)	(0.0229)	(0.0167)	(0.0288)	(0.0328)
	β_1	0.269^a	0.103^a	0.070^a	0.157^a	0.074^a	0.059^a	0.250^a	0.139^a	0.085^a
		(0.0313)	(0.0195)	(0.0162)	(0.0316)	(0.0196)	(0.0162)	(0.0307)	(0.0239)	(0.0217)
	β_2	0.273^a	0.266^a	0.271^a	0.452^a	0.385^a	0.328^a	0.397^a	0.357^a	0.356^a
		(0.0518)	(0.0531)	(0.0499)	(0.0559)	(0.0542)	(0.0508)	(0.0535)	(0.0607)	(0.0665)
	β_3	0.193^a	0.236^a	0.183^a	0.125^a	0.206^c	0.190^a	0.320^a	0.386^a	0.397^a
		(0.0531)	(0.0634)	(0.0601)	(0.0636)	(0.0718)	(0.0683)	(0.0593)	(0.0746)	(0.0754)
MIDAS	μ_H	0.833^a	1.608^a	1.932^a	0.887^a	1.707^a	2.007^a	1.075^a	1.911^a	2.216^a
		(0.1100)	(0.1695)	(0.1957)	(0.1156)	(0.1213)	(0.1364)	(0.0416)	(0.0619)	(0.0766)
	ϕ_H	0.635^a	0.485^a	0.429^b	0.695^a	0.578^a	0.498^a	0.906^a	0.803^a	0.733^a
		(0.1042)	(0.1700)	(0.1991)	(0.1200)	(0.1177)	(0.1326)	(0.0337)	(0.0457)	(0.0562)
	θ_1	-0.107^a	−0.084	−0.075	-0.088^a	−0.039	−0.039	-0.104^a	-0.059^a	-0.053^a
		(0.0241)	(0.0654)	(0.0954)	(0.0329)	(0.0235)	(0.0304)	(0.0150)	(0.0123)	(0.0142)
	θ_2	0.001^a	0.001	0.001	0.001^b	0.000	0.000	0.001^a	0.001^a	0.001^a
		(0.0002)	(0.0008)	(0.012)	(0.0004)	(0.0003)	(0.0004)	(0.0002)	(0.0001)	(0.0002)

Note: The dependent variable for all forecast horizons H is the logarithm of the square root of realized volatility. The regressors are based on squared returns. Shown in parentheses are Newey-West HAC standard errors. The coefficient estimates significant at 1, 5 and 10 percent are denoted with superscripts a, b, and c, respectively. The in-sample period runs from January 3, 2000 to January 25, 2007 (1,750 obs.)