JEL Classification: C1, C5, G1
Keywords: intraday data, realized variance, return and volatility distributions, heterogeneous autoregressive model

Distribution and Dynamics of Central-European Exchange Rates: Evidence from Intraday Data*

Vít BUBÁK – Institute of Economic Studies, Charles University in Prague and CES, Université Paris I. Panthéon-Sorbonne (vit.bubak@malix.univ-paris1.fr)
Filip ŽIKEŠ – Imperial College London, Business School (fzikes@imperial.ac.uk)

Abstract
This paper investigates the behavior of the EUR/CZK, EUR/HUF and EUR/PLN spot exchange rates in the period 2002–2008, using 5-minute intraday data. We find that daily returns on the corresponding exchange rates scaled by model-free estimates of daily realized volatility are approximately normally distributed and independent over time. On the other hand, daily realized variances exhibit substantial positive skewness and very persistent, long-memory type of dynamics. We estimate a simple three-equation model for daily returns, realized variance and the time-varying volatility of realized variance. The model captures all salient features of the data very well and can be successfully employed for constructing point, as well as density forecasts for future volatility. We also discuss some issues associated with measuring volatility from the noisy high-frequency data and employ a simple correction that accounts for the distortions present in our dataset.

1. Introduction

The recent economic downturn has put an end to a period of relative stability that the Czech koruna, Hungarian forint and Polish złoty enjoyed over the last years. The considerable increase in the volatility of these currencies raises a question about the ability of the Czech Republic, Hungary and Poland to fulfill the exchange rate stability criteria stipulated in the Maastricht Treaty. Indeed, these criteria require that for at least two years prior to the entry into the Eurozone, the applicant country's currency remain within a normal fluctuation band around the central parity, effectively setting limits to the currency's volatility during the pre-accession period (Antal and Holub, 2007). There is no doubt that while the choice of the appropriate monetary and exchange rate policies will be crucial to ensure that the currency meets the convergence criteria, the design and implementation of such policies would not be possible without a thorough understanding of the statistical properties of the currencies in question. A practical framework for accurate modeling and forecasting of the exchange rate volatility in particular could ultimately help in making the relevant policies more efficient.

A good knowledge of the Central European (CE) exchange rates dynamics is equally relevant for asset pricing and risk management. Understanding the conditional probability distribution of the exchange rate returns and their volatility is critical for accurate estimation of various models used in pricing and hedging derivative securities written on the exchange rate. On a more general level, frequent and poten-

* The authors thank Jozef Baruník for helpful comments on an earlier version of this paper. Bubak gratefully acknowledges financial support from the Czech Ministry of Education, grant code MSM0021620841.

tially large unexpected exchange rate movements adversely affect the performance of export-oriented businesses. Papaioannou (2006) discusses specific types of exchange rates risk that these companies face at times of increased currency volatility, includeing transaction costs associated with hedging against unfavorable exchange rate movements and economic costs arising from increased uncertainty about future relative competitiveness. As the CE currencies continue to suffer from relatively high volatility triggered by the global economic crisis, containing these and related risks demands effective risk management decisions that are impossible without a sound knowledge of the underlying exchange rate behavior.

The CE currencies have been subject to a wide range of studies. The most recent focus on understanding the effectiveness of foreign exchange interventions conducted by Central Banks (see Geršl, 2004; Geršl, 2006; Geršl and Holub, 2006; Égert and Komárek, 2006; Égert, 2007), the sustainability of the real exchange rates (Bulíř and Šmídková, 2005), or the equilibrium real exchange rate determination (Melecký and Komárek, 2008), among others.

In contrast, only a limited number of studies have attempted to model the dynamics of the spot exchange rates for the CE currencies. Kočenda and Valachy (2006) provide a detailed analysis of the exchange rate volatility in the Visegrád countries, with a particular focus on the period in which these countries abandoned tight FX regimes for more flexible ones. Using daily nominal exchange rate data, the authors employ an augmented version of a threshold GARCH-in-Mean (T-GARCH) model to study the effects of path dependency, asymmetric shocks, and movements in interest rates on exchange rate volatility during the regime transition. The study shows that the introduction of the more flexible regime lead to a general increase in exchange rate volatility, with the level of volatility persistence becoming roughly the same across the exchange rates analyzed. The authors also find a significant and negative effect of asymmetric shocks on the volatility of Polish złoty and Hungarian forint under the floating regime.

In a related paper, Fidrmuc and Horváth (2008) analyze the exchange rate dynamics in the selected EU members including the Czech Republic, Hungary and Poland, using daily data from 1999 to 2006. The authors apply both a GARCH model and an extended version of the TARCH model to assess the exchange rate volatility in connection with the estimated target exchange rate and the credibility of exchange rate management. Among other findings, the study shows that the daily exchange rate volatility exhibits strong persistence as well as systematic asymmetric effects, with the latter being especially pronounced during the periods of exchange rate appreciation.

Horváth (2005) investigates the medium-term determinants of the bilateral exchange rate volatility of Central and Eastern European countries (CEEc) based on the optimal currency area criteria. As part of the analysis, the author also compares the actual and predicted exchange rate variability between the Euro area countries and the CEEc. Although limited to the use of quarterly data and a relatively short sample period from 1999 to 2004, the study shows that the actual exchange rate variability is larger in the CEEc compared to what it had been in the Euro area before its creation. In addition, the author finds the predicted exchange rate variability to be close to the Eurozone levels, with the difference between the latter and the actual variability caused by the Euro area countries participating in the ERM during the sample period.

Finally, Frömmel (2007) provides an interesting investigation of the changes between volatility regimes in five Central and Eastern European countries, including the Czech Republic, Poland and Hungary. Frömmel employs a Markov-Switching GARCH model to study whether the changes between the volatility regimes are consistent with changes in the official exchange rate arrangements. Among other findings, the author concludes that an increase in the flexibility of the exchange rate regime leads to an increase in exchange rate volatility.

The goal of this paper is to examine the conditional distribution of the Czech koruna, Hungarian forint and Polish złoty exchange rates vis-à-vis the Euro in the period 2002–2008. Employing a 5-minute intraday data, we examine the distributional properties and time-series dynamics of both daily exchange rate returns, as well as daily realized variance. Unlike the existing empirical literature that employs almost exclusively a GARCH framework to study the dynamics of the exchange rate, our work relies on model-free nonparametric measures of ex-post volatility based on the use of intraday data. This approach, pioneered by Andersen and Bollerslev (1998), has attracted substantial attention in the recent financial econometric literature; see e.g. McAleer and Medeiros (2008) for a recent review. It offers a number of advantages.

First, no parametric assumptions are needed to ensure that the realized variance and related measures are consistent for the true, unobserved volatility, apart from some mild regularity conditions. This is in stark contrast to the GARCH framework, where all results concerning the behavior of volatility hinge on a particular specification of the GARCH variance equation.

Second, realized variance captures the total variation in the price or exchange rate over a given period of time, unlike a GARCH-type model that focuses on conditional volatility of the price at time t, given the information set available at time $t - 1$. In other words, realized variance combines both the volatility expectations as well as the innovations to volatility. This carries important implications for studying the conditional distributions of one-period returns as pointed out by Andersen, Bollerslev and Dobrev (2007): while the one-period financial returns standardized by conditional volatility typically appear to be leptokurtic, standardizing by realized volatility produces approximately Gaussian innovations. This in turn lends empirical support to a large class of continuous-time stochastic volatility models widely employed in the asset pricing literature.

Finally, since the realized variance and alternative related measures render volatility essentially observable up to a measurement error that vanishes as the sampling frequency increases, simple time-series models can be used to model and accurately forecast future volatility (see Andersen, Bollerslev, Diebold and Labys, 2003; Andersen, Bollerslev and Dobrev, 2007, among others). This includes not only point forecasts, that is, the expected future volatility, but the entire predictive density for future volatility, allowing for construction of confidence intervals around the point forecast or, similarly, estimation of the probability that future volatility remains within a certain fluctuation band. The ability to provide the predictive density for future volatility also facilitates the measurement and management of risk associated with trading realized volatility, which has become very popular in recent years (e.g. Bondarenko, 2007). In this paper, we only focus on a simple model for returns and variance since our primary interest lies in studying the dynamics and conditional distributions of the EUR/CZK, EUR/HUF and EUR/PLN spot exchange rates.

Our empirical results confirm some stylized facts about the behavior of returns and volatility of foreign exchange rates. We find that daily returns on the exchange rates are approximately normally distributed and independent over time, when properly scaled by model-free estimates of daily realized variance. Daily realized variance, on the other hand, exhibits substantial positive skewness as well as a very persistent, long-memory type of dynamics. We propose a relatively simple model for daily returns, realized variance and the time-varying volatility of realized variance, finding that it very well captures all salient features of the data. In addition, the model is shown to perform remarkably well out-of-sample, delivering accurate volatility forecasts. It may therefore serve well as an auxiliary model for estimating various continuous-time stochastic volatility models used for pricing derivative securities written on the exchange rate (Bollerslev, Kretschmer, Pigorsch and Tauchen, 2009).

The rest of the paper is organized as follows. In Section 2 we describe our theoretical framework and discuss some distributional predictions that it generates for the EUR/CZK, EUR/HUF and EUR/PLN returns. In Section 3, we follow with a definition of the realized variance as a model-free measure of variation in asset prices and some of the issues associated with measuring volatility from noisy high--frequency data. In Section 4 we describe the data and in Section 5 we report the empirical results. In particular, we present the results of the tests of normality and independence of returns standardized by realized volatility, the estimation of a joint model for daily returns, realized variance and the volatility of realized variance, and the results of an out-of-sample volatility forecasting exercise. Section 6 concludes the paper with some suggestions for future work.

2. Theoretical Framework

Following a vast body of recent literature in financial econometrics, we adopt a relatively simple, yet very general continuous-time framework. Working in continuous time has a number of technical advantages, but more importantly it provides a direct link to the asset pricing literature, which establishes a number of important results concerning the restrictions on admissible models governing asset prices in an arbitrage-free environment (Back, 1991). A detailed overview of this and related issues is beyond the scope of this paper and we refer the interested reader to an excellent discussion in Andersen, Bollerslev, Diebold and Labys (2003).

We assume that the logarithmic spot exchange rate, s_t, follows a stochastic volatility model given by

$$s_t = \int_0^t \mu_u \mathrm{d}u + \int_0^t \sigma_u \mathrm{d}W_u \tag{1}$$

where μ_t and σ_t denote the drift and volatility processes, respectively, and W_t is a standard Brownian motion. Both μ_t and σ_t are allowed to be general stochastic processes and we do not impose any parametric assumption regarding their respective laws of motion. Also, no restrictions are placed on the dependence between volatility (σ_t) and the Brownian motion (W_t) driving the exchange rate innovations.

A few remarks regarding the model in equation (1) are in order. First, the sample paths of the exchange rate are continuous, hence ruling out the presence of jumps. We choose to make this assumption to keep our framework simple for the sake of exposition, but nothing prevents us from including a jump process to the drift and diffusion components in equation (1). Indeed, the measures of volatility that we employ later in the paper can capture both parts of the variation, i.e. the diffusion part and jump part, if present, and hence there is no loss of generality in this sense by doing otherwise.

Second, the model nests a wide variety of arbitrage-free stochastic volatility models employed in the asset pricing literature. The well-known Black-Scholes model, where both the drift and volatility are constant, is a prominent example. For more general and empirically relevant specification see Chernov, Gallant, Ghysels and Tauchen (2003) and the references therein.

Finally, the model delivers testable distributional predictions: the one-period returns defined as $r_t = s_t - s_{t-1}$ are conditionally on the sample path of drift and volatility, normally distributed. Formally:

$$r_t \mid \{\mu_u, \sigma_u\}_{t-1}^{t} \sim N\left(\int_{t-1}^{t} \mu_u \mathrm{d}u, \int_{t-1}^{t} \sigma_u^2 \mathrm{d}u\right) \tag{2}$$

Since the drift is typically negligible at daily and weekly frequencies, especially in the case of foreign exchange rates, the key quantity that we are interested in is the so-called integrated variance,

$$IV_t = \int_{t-1}^{t} \sigma_u^2 \mathrm{d}u \tag{3}$$

which, as equation (2) shows, is the natural measure of variation in the one-period returns. The conditional normality of r_t further implies that in the absence of dependence between the volatility process and the Brownian motion driving the exchange rate (W_t), the one-period standardized returns follow the standard normal distribution,

$$\frac{r_t - \int_{t-1}^{t} \mu_u \mathrm{d}u}{\left(\int_{t-1}^{t} \sigma_u^2 \mathrm{d}u\right)^{1/2}} \overset{iid}{\sim} N(0,1) \tag{4}$$

Similar predictions can be derived when the volatility process correlates with the Brownian motion. The normality of properly standardized returns has found an overwhelming empirical support across different assets classes; see e.g. Andersen, Bollerslev and Dobrev (2007), Andersen, Bollerslev, Frederiksen and Nielsen (2009) and Žikeš (2008) for recent evidence from equity index futures, individual stocks, and foreign exchange rates, respectively. It is worth reiterating that this distributional assumption can be tested without making any parametric assumptions about the volatility process since the integrated volatility appearing in the denominator of the standardized returns can be consistently estimated by nonparametric methods, which we describe in the next section.

3. Measuring Daily Variance

Suppose we obtain a sample of size $T(M+1)$ corresponding to T days, each having $M + 1$ intraday observations of the logarithmic spot exchange rate. We denote by

$s_{t(i)}$ the i-th observation of the log-spot rate on day t, with $i = 0,\ldots, M$ and $t = 1,\ldots,T$. The simplest and most widely used estimator, the well-known realized variance[1] (Andersen and Bollerslev, 1998), is obtained by summing the squared intraday returns:

$$RV_{t,M} = \sum_{i=1}^{M} (\Delta_i s_t)^2 \tag{5}$$

where $\Delta_i s_t = s_{t(i)} - s_{t(i-1)}$ denotes the i-th intraday return on day t. As the sampling frequency increases, $M \to \infty$, the realized variance converges in probability to the integrated variance, IV_t (see e.g., Protter, 2005). Moreover, under some mild regularity conditions, a central limit theorem can be obtained, establishing the $M^{1/2}$ rate of convergence (Barndorff-Nielsen and Shephard, 2002). Thus, the realized variance is a fully nonparametric estimator of the integrated variance, yet it achieves the usual parametric rate of convergence. We finally remark that if jumps are present in the true price process in addition to the diffusion component, the realized variance will pick up both: it will converge to the integrated variance plus the sum of squared jumps, thereby providing a measure of the overall variation in the one-period returns.

The asymptotic results mentioned above seem to suggest that one should sample as frequently as possible to achieve highly accurate realized variance estimates. However, when taken to the data, one quickly realizes that this is actually not optimal. The reason is that intraday data sampled at very high frequencies tend to be contaminated by the so-called microstructure noise. The noise arises from a number of frictions inherent to the process of trading and posting bid and ask quotes. See O'Hara (1995) for an overview of the theory market microstructure and Hansen and Lunde (2006) for the implications of the presence of noise for estimating volatility from high-frequency data.

A typical approach to modeling the noise in the realized variance literature is to assume that the noise is additive, i.e.

$$s^*_{t(i)} = s_{t(i)} + \varepsilon_{t(i)}\,,\ \ \varepsilon_{t(i)} \sim D(0,\omega^2) \tag{6}$$

where $s_{t(i)}{}^*$ is the actual price the econometrician observes, while the efficient price, $s_{t(i)}$, remains unobserved due to contamination by $\varepsilon_{t(i)}$. Earlier contributions assumed that the noise is independently and identically distributed over time and is independent from the efficient prices. Both have been gradually relaxed and the estimator we use in this paper works under very general conditions. Nonetheless, the assumption of *i.i.d.* turns out to be approximately satisfied in foreign exchange data and as we will see below, also for the exchange rates analyzed in the current study. We will therefore retain this assumption for the sake of exposition.

An immediate consequence of the presence of *i.i.d.* microstructure noise is that the realized variance becomes biased and inconsistent as the sampling frequency increases. The noise contaminating the efficient price induces a moving-average type of structure in the observed intraday returns,

$$\Delta_i s^*_t = \Delta_i s_t + \varepsilon_{t(i)} - \varepsilon_{t(i-1)} \tag{7}$$

[1] Note that it is common in the literature to abuse terminology by using 'realized variance' and 'realized volatility' interchangeably to refer to the same quantity defined in equation (5). We will try to avoid this by reserving the term 'realized variance' for $RV_{t,M}$ defined as in (5) and 'realized volatility' for $(RV_{t,M})^{1/2}$.

As a result, the realized variance behaves, for large M, as

$$RV_{t,M} \approx IV_t + 2M\omega^2 \tag{8}$$

and is thus biased and inconsistent as it tends to infinity with $M \to \infty$. In fact, for large M, the realized variance, when scaled by $2M$ can be used to estimate the variance of the noise (Bandi and Russell, 2006).

The vast majority of papers in the literature circumvent the problem of noise by sampling sparsely, that is, by sampling at frequencies at which the bias is small. To this end, Andersen, Bollerslev, Diebold and Labys (2000) introduce the so-called volatility signature plot that shows the average daily realized volatility calculated at different sampling frequencies. In the absence of noise, this plot should be flat. On the other hand, if the noise is present, the signature plot will reveal the frequency at which the bias induced by it kicks in. This frequency is then used in the empirical work to measure the daily volatility.

Sampling sparsely, however, entails throwing away a lot of data, which violates one of the main rules of statistics (Zhang, Mykland and Aït-Sahalia, 2005). Therefore a number of solutions have been proposed in the literature to correct the biases associated with microstructure noise directly. Here we use the moving-average based estimator first used by Andersen, Bollerslev, Diebold and Labys (2001) and recently theoretically studied by Hansen, Large and Lunde (2008), since the microstructure noise contaminating our data seems to exhibit a simple *i.i.d.* structure. For an alternative approach, see Barndorff, Nielsen, Hansen, Lunde and Shephard (2008) and the references therein.

The moving-average based estimator exploits the MA(1) structure of observed returns, $\Delta_i s_t^*$. The intraday returns are first filtered by an MA(1) model,

$$\Delta_i s_t^* = \eta_{t,i} - \theta \eta_{t(i-1)} \tag{9}$$

where the parameter θ can be estimated by the method of quasi maximum likelihood. In the second step, the usual realized variance is applied to the filtered intraday returns, $\hat{\eta}_{t(i)}$, i.e.

$$RVMA_{t,M}^{(1)} = (1-\hat{\theta})^2 \sum_{i-1}^{M} \hat{\eta}_{t(i)}^2 \tag{10}$$

where the scaling constant $(1-\hat{\theta})^2$ ensures that the estimator is unbiased and consistent for the integrated variance. Hansen, Large and Lunde (2008) provide a central limit theorem for $RVMA_{t,M}^{(1)}$ and establish the $M^{1/4}$ rate of convergence, which is known to be the best possible rate when estimating volatility from noisy data.

The discussion of the issues associated with measuring volatility from noisy high-frequency data completes the methodology part of the paper. Before we turn to the empirical application we first carefully describe the data.

4. Data Description and Preliminaries

We employ 5-minute spot exchange rate mid-quotes covering the period from January 4, 2002 through December 30, 2008 for the case of EUR/CZK and EUR/PLN and from January 2, 2003 through December 30, 2008 for the case of EUR/HUF.

Figure 1

Left: Volatility signature plots for the *RV* (defined in (5)) and the *RVMA* (lower line). The numbers on the *x*-axis correspond to 5, 10, 15, 20, 30, and 60 min sampling frequency. *Right:* ACF of intra-day returns obtained by tick-time sampling. The first (second, third) row corresponds to the case of EUR/CZK (EUR/HUF, EUR/PLN) 5 min exchange rate returns, respectively.

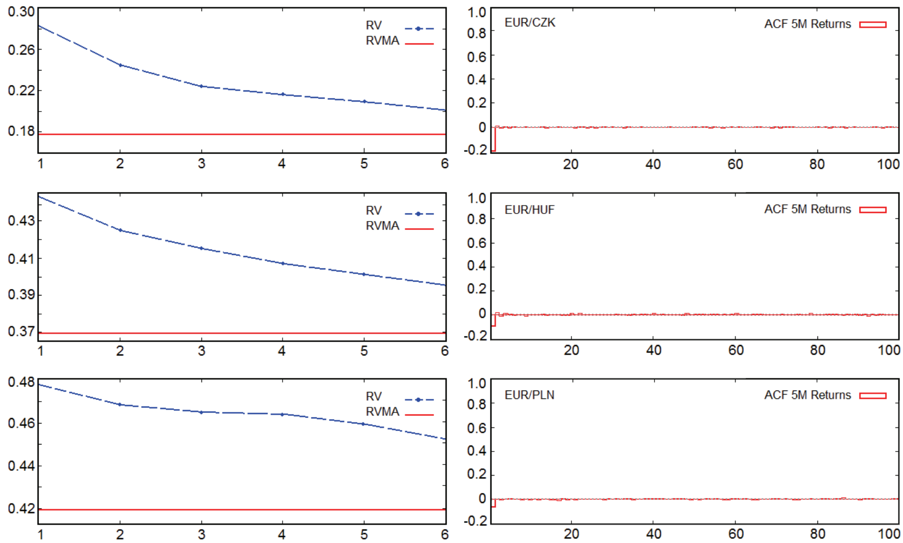

The mid-quotes are constructed by taking the average of the best bid and ask quotes available at the end of each 5-minute interval. The data was obtained from Olsen Financial Technologies.

Similarly to other FX markets, the EUR/CZK, EUR/HUF and EUR/PLN markets operate 24 hours per day. To avoid distortions associated with illiquidity and thin trading, we follow the usual approach in the literature and discard weekend periods from Friday 21:00 GMT (22:00 CET) until Sunday 21:00 GMT, as well as holidays. This leaves us with a total of 1,780, 1,507, and 1,762 trading days for the EUR/CZK, EUR/HUF and EUR/PLN exchange rates, respectively.

We define a trading day on the interval from 21:00 GMT to 21:00 GMT of the following day as is common in the literature (e.g. Andersen, Bollerslev, Diebold and Labys, 2003). Since the trading activity in the foreign exchange markets exhibits substantial deterministic intraday variation, we resort to tick-time sampling for the purposes of measuring daily volatility (e.g. Oomen, 2006). That is, for each day in the sample, the series of intraday prices are obtained by discarding duplicate quotes, and the intraday returns are then calculated from these generally irregularly-spaced prices. This procedure eliminates the zero intraday returns largely prevalent in periods of thin trading and makes the resulting irregularly-spaced intraday returns closer to being homoskedastic. Theory implies that this should generally improve the accuracy of volatility estimation (e.g. Oomen, 2006), and is particularly desirable when employing the moving-average based estimator (Hansen, Large and Lunde, 2008).

We now proceed to discuss the problem of measuring daily volatility. The autocorrelation functions (ACF) of the intraday returns, plotted in *Figure 1* (left column),

Figure 2
Plots of daily spot rates, *left*, and of daily returns, *right*, for the case of EUR/CZK (first row), EUR/ /HUF (second row) and EUR/PLN (third row) exchange rates

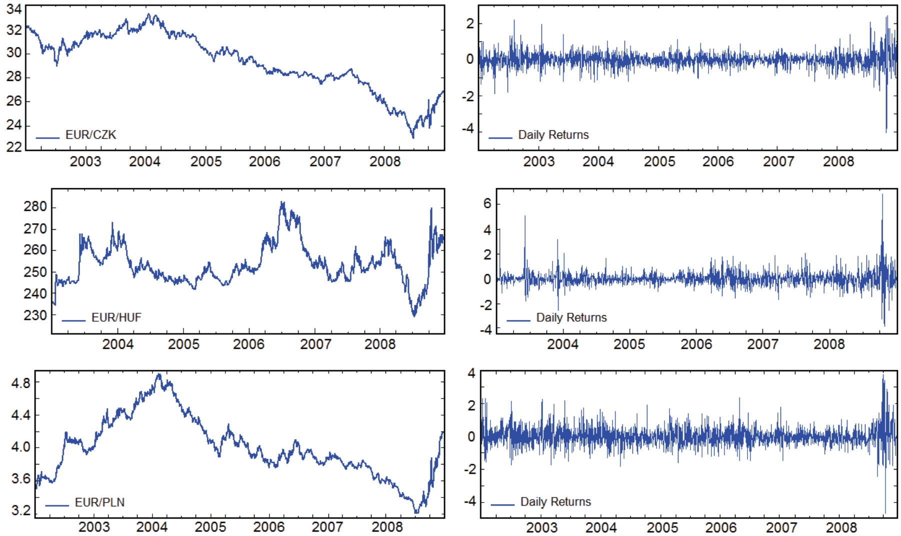

all exhibit a significant negative spike at lag one, with essentially no statistically significant autocorrelation at longer lags. This is consistent with the intraday returns having an MA(1) component induced by an *i.i.d.* microstructure noise contaminating the spot exchange rate. The usual realized variance estimator will be therefore substantially upward biased at the 5-minute and perhaps even at lower sampling frequencies. Given the simple dynamics of the noise process implied by the ACF, the bias should be to a large extent corrected by the moving-average based estimator, $RVMA_{t,M}$.

To see this, we plot in *Figure 1* (right column) the realized variance signature along with the average $RVMA_{t,M}$. The bias of $RV_{t,M}$, increasing in the sampling frequency is clearly apparent from the plot. Even at frequencies as low as 30 minutes the usual realized variance still exhibits a large positive bias. In the rest of the paper, we therefore employ the moving-average based estimator as our preferred measure of volatility and to simplify the notation, we reserve RV_t to denote $RVMA_{t,M}$ since no confusion should arise regarding the particular realized variance measure and sampling frequency used.

The plots of EUR/CZK, EUR/HUF and EUR/PLN spot exchange rates and the corresponding daily returns are presented in *Figure 2*. The reader will immediately notice several periods of increased volatility that characterize the daily exchange rate returns. In particular, a significantly larger volatility around the last 120 days of the samples that mirrors a sharp depreciation of the currencies in the last quarter of 2008, as well as an overall increase in uncertainty associated with the global economic downturn is common to all three currencies. Still, other periods of larger volatility can be discerned. For example, an increased exchange rate volatility is evident during the first 250 days of the sample of EUR/CKZ returns that reflects an episode of relatively strong nominal appreciation of the Czech currency that started in 2001 and con-

Table 1 Descriptive Statistics

Descriptive statistics for daily returns, daily realized variance, daily logarithmic realized variance and daily returns standardized by realized volatility. The realized variance is calculated using the moving-average estimator. The sample runs from January 4, 2002 to December 30, 2008 for EUR/CZK and EUR/PLN, and from January 2, 2003 to December 30, 2008 for EUR/HUF

		Mean	Std Dev	Skew	Kurt	Min	Max
CZK	r_t	-0.010	0.435	-0.389	13.30	-4.058	2.422
	RV_t	0.179	0.287	7.309	85.81	0.010	5.139
	$\log(RV_t)$	-2.166	0.826	0.845	-4.246	-4.640	1.637
	$r_t/(RV_t)^{1/2}$	-0.034	0.975	-0.046	3.028	-4.046	2.957
HUF	r_t	0.008	0.618	1.700	22.45	-3.797	6.850
	RV_t	0.370	0.816	10.64	162.5	0.011	16.75
	$\log(RV_t)$	-1.594	0.967	0.587	4.029	-4.529	2.818
	$r_t/(RV_t)^{1/2}$	-0.037	0.913	0.113	2.952	-2.769	3.012
PLN	r_t	0.009	0.606	0.335	8.946	-4.640	3.747
	RV_t	0.419	0.679	6.445	60.57	0.025	10.25
	$\log(RV_t)$	-1.324	0.847	0.671	4.295	-3.705	2.327
	$r_t/(RV_t)^{1/2}$	-0.041	0.946	-0.007	2.771	-3.110	3.356

tinued throughout 2002 and was driven mainly by market expectations of significant Euro-denominated privatization revenues being converted into the domestic currency (see Geršl (2004) for a full description of the underlying events). Similarly, strong downward pressures on Hungarian forint due to weak economic outcomes and market doubts about the consistency of monetary policies resulted in a larger EUR/HUF volatility during the second half of 2003.

5. Empirical Results

5.1 Distribution of Daily Returns and Realized Variance

We begin our analysis by looking at the properties of the *raw* daily returns. The statistics reported in *Table 1* indicate that the daily returns exhibit excess kurtosis and are either slightly negatively (EUR/CZK) or positively skewed (EUR/HUF, EUR/PLN), relative to normal distribution. The EUR/HUF returns exhibit the largest degree of skewness and excess kurtosis among the three exchange rates.

To investigate the distributional properties of daily returns standardized by realized volatility, $r_t/(RV_t)^{1/2}$ (also plotted in the right column of *Figure 5*), we run a battery of tests. Recall that (4) implies a sharp null hypothesis of standard normality and independence. Thus we first consider the moment-based test of Bontemps and Meddahi (2005), focusing on the first four moments of standardized returns. Under the null hypothesis, the standardized returns have zero mean, unit variance, zero skewness and kurtosis equal to three. We focus on two versions of the test $H_{1\text{–}4}$, which has as its null hypothesis that all four moments are equal to those of a standard normal, and $H_{3\text{–}4}$, which only takes into account the third and fourth moment. The latter is asymptotically equivalent to the well-known Jarque-Bera test. $H_{1\text{–}4}$ is asymptotically distributed as a $\chi^2(4)$ random variable, while $H_{3\text{–}4}$ as a $\chi^2(2)$.

Second, we employ the well-known Kolmogorov-Smirnov (KS) test for the null hypothesis of standard normality. Unlike the moment-based test, the KS test is a consistent test, i.e. it has asymptotically unit power against all alternatives. The limiting distribution of the KS test statistic is non-standard and the critical values have to be simulated.

Finally, we run the test developed by Hong and Li (2005) (HL) to test the null hypothesis of both standard normality and independence. Since this test is not widely used a brief description is in order. The test is based on the observation that under the null, the joint distribution of $r_t/(RV_t)^{1/2}$ and $r_{t-k}/(RV_{t-k})^{1/2}$ factorizes into the product of two standard normal marginals for any k. Hong and Li (2005) propose to estimate the joint distribution by nonparametric methods and build a test statistic based on the integrated squared difference between the estimated joint density and the joint density under the null hypothesis. If the null hypothesis is true, the difference should be small. Under the alternative, the test statistic diverges. The limiting distribution of the HL test statistic is standard normal for any k and the test statistics are asymptotically independent across different k's. A joint χ^2 test can be therefore easily constructed by taking the sum of squared HL test statistics for different k's.

Before we turn to the empirical results, it is worth mentioning that testing for normality and independence of returns standardized by realized volatility or any other consistent measure of integrated variance entails a measurement error problem. The null hypothesis is specified in terms of the unobserved standardized returns, that is, by the returns scaled by the true integrated volatility. Replacing the unobserved volatility by its sample counterpart induces a measurement error that may in turn affect the central limit of the test statistics. Žikeš (2008) recently studies this problem for the tests described above and establishes primitive conditions on the spot volatility process as well as the restrictions on the relative rate of growth of T and M such that the measurement error vanishes asymptotically. For finite T and M, important distortions may arise and this has to be kept in mind when interpreting the results of the tests. In other words, a rejection of the null hypothesis with relatively small M may indeed be a symptom of the measurement error, rather than a genuine feature of the data.

Table 2 summarizes the results of the normality tests. Consistent with the prediction of the model in equation (4), we find that the sharp null hypothesis of standard normality is not rejected by any of the tests at conventional significance levels for EUR/CZK. Both the moment-based test statistics as well as the Kolmogorov-Smirnov test statistics are well below their respective critical values. The excellent fit of the standard normal distribution for the standardized returns is also apparent from the relevant kernel density plot reported in *Figure 3* (right column) and the last column of *Table 1*. The estimated density is essentially indistinguishable from $N(0,1)$, with the mean, standard deviation, skewness and kurtosis appearing very close to those of a standard normal distribution.

The HL test fails to detect any dependence in the time series of the standardized residuals up to lag 5. The p-value corresponding to lag 2 is close to 5%, which may suggest some dependence between $r_t/(RV_t)^{1/2}$ and $r_{t-2}/(RV_{t-2})^{1/2}$. However, given that we run the test for a number of lags, the relevant test statistics to look at is the one for the joint test HL_{joint}, which has a p-value of 0.425 and thus clearly fails to reject the null hypothesis of standard normality and independence. The conclusion of the HL test is further corroborated by the autocorrelation function for the standard-

Table 2 Tests for Normality and Independence of Standardized Returns

Left panel reports tests for standardized returns, $r_t/(RV_t)^{1/2}$, while the *right* panel reports tests for standardized returns that were studentized by sample mean and standard deviation. $H_{1–4}$ denotes a test statistic for the null hypothesis of standard normality based on the first four Hermite polynomials. Similarly, $H_{3–4}$ denotes a test statistic for the null hypothesis of normality based on the third and fourth Hermite polynomials. *KS* denotes the Kolmogorov-Smirnov test statistic for the null hypothesis of standard normality and HL_k the test statistic for the null hypothesis of both standard normality and independence in the standardized returns at lag *k* or a joint test. *P*-values are reported in parentheses, except for the KS test, where we report the 5% critical values instead. We denote by * the test statistic that exceeds its 5% critical value.

	A. Std. Returns			B. Studentized Std. Returns		
	CZK	HUF	PLN	CZK	HUF	PLN
$H_{1–4}$	4.682 (0.321)	25.50* (0.000)	14.01* (0.007)	–	–	–
$H_{3–4}$	0.521 (0.771)	2.821 (0.244)	1.683 (0.431)	0.679 (0.712)	3.339 (0.188)	3.851 (0.146)
KS	0.982 (1.36)	1.853* (1.36)	1.050 (1.36)	0.445 (0.89)	0.647 (0.89)	0.583 (0.89)
HL_1	1.183 (0.118)	3.538* (0.000)	0.879 (0.190)	0.494 (0.311)	-0.017 (0.507)	0.180 (0.429)
HL_2	1.611 (0.054)	3.471* (0.000)	0.977 (0.164)	0.848 (0.198)	0.073 (0.471)	0.048 (0.481)
HL_3	-0.132 (0.553)	2.607* (0.005)	-0.785 (0.784)	-0.907 (0.818)	-0.766 (0.778)	-1.819 (0.966)
HL_4	0.874 (0.191)	3.006* (0.001)	0.432 (0.333)	0.104 (0.458)	-0.460 (0.677)	-0.503 (0.693)
HL_5	1.195 (0.116)	3.362* (0.000)	1.794* (0.036)	0.483 (0.315)	0.204 (0.419)	0.757 (0.224)
HL_{joint}	6.201 (0.287)	51.70* (0.000)	5.748 (0.332)	2.029 (0.845)	0.845 (0.974)	4.172 (0.525)

ized returns plotted in *Figure 4* (right column). All autocorrelation coefficients plotted there remain within the confidence bands, indicating no omitted dynamics.

Turning to the standardized returns of EUR/HUF and EUR/PLN we find that the test based on the first four Hermite polynomials $H_{1–4}$, rejects the null hypothesis of standard normality. In case of EUR/HUF, the same conclusion is obtained by the other tests as well, except for the test based on the third and fourth Hermite polynomials ($H_{3–4}$). This suggests that the rejection of the null hypothesis may be due to the mean and/or standard deviations of the standardized returns deviating from zero and one, respectively. Looking at the descriptive statistics reported in *Table 1* we indeed observe that the standard deviation of $r_t/(RV_t)^{1/2}$ is smaller than one, while the mean appears to be indistinguishable from zero.

To see if the departure of the standard deviation from one is responsible for the rejection of the sharp null hypothesis of standard normality, we next run the normality test on studentized standardized returns obtained by de-meaning and dividing the original standardized returns by their sample standard deviation. This of course introduces a parameter uncertainty problem since we do not know the true mean and standard deviation. Fortunately, all tests but the KS test employed here are robust to this problem and hence valid inference is obtained by replacing the true parameters by consistent estimates. In case of the KS test, we use the Lilliefors approximate crit-

Figure 3 Estimated Kernel Densities for the Daily Logarithmic Realized Variance (log (*RV*)), *left*, and Daily Standardized Returns ($r_t/(RV)^{1/2}$), *right*
constructed using EUR/CZK (first row), EUR/HUF (second row), and EUR/PLN exchange rate returns (third row)

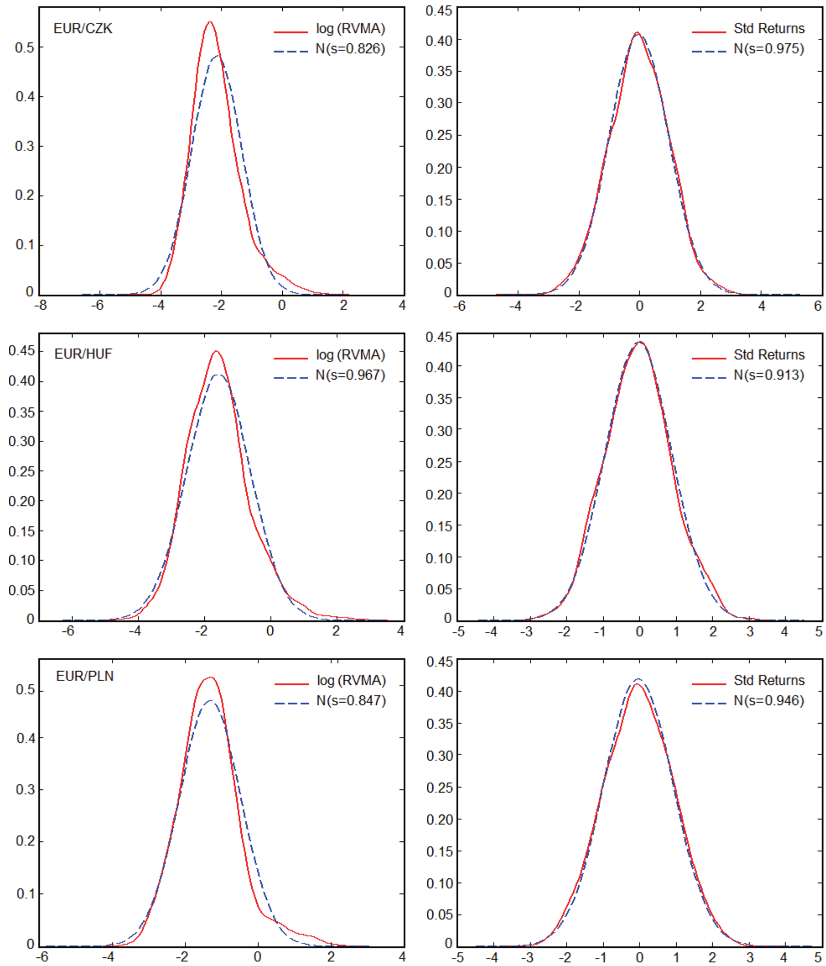

ical value. The results are reported in Panel B of *Table 2*. Clearly, the null hypothesis of normality and independence is not rejected by any test, confirming our initial conjecture regarding the departure from the sharp null of standard normality.

Turning to realized variance, *Figure 3* (left column) presents the nonparametric estimates of the density of logarithmic realized variance for the three exchange rates, while *Table 1* provides additional information. We observe that even after taking the logarithmic transformation the distribution of realized variance exhibits positive skewness. This has important implications for modeling and forecasting the distribution of future volatility and will be explicitly taken into account when constructing a joint model for returns and volatility in the next section of the paper. To test the null hypo-

Figure 4 ACF Plots of Daily Realized Variance (*RV*), *left*, and of Daily Standardized Returns ($r_t/(RV)^{1/2}$), *right*
for the case of EUR/CZK (first row), EUR/HUF (second row), and EUR/PLN exchange rate returns (third row)

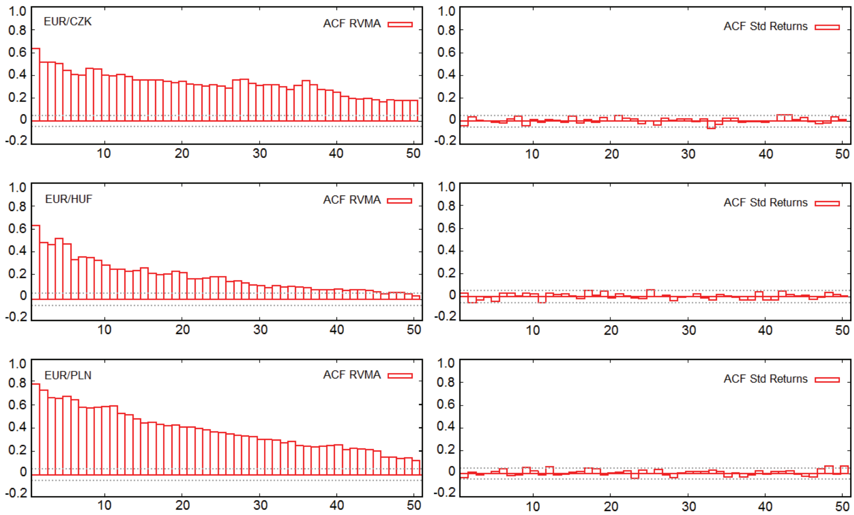

thesis of normality of logarithmic realized variance formally, we employ the test based on the third and fourth Hermite polynomials ($H_{3\text{–}4}$) constructed using the Newey-West weighting matrix. This test is valid in the presence of parameter uncertainty as well as dependence in the logarithmic realized variance, unlike the other tests we used before. The test statistics read 103.8, 56.6 and 50.3 for EUR/CZK, EUR/HUF and EUR/PLN, respectively, clearly rejecting the null hypothesis. Therefore neither the realized variance nor its logarithmic transformation follows the normal distribution.

Finally, we examine the dynamics of realized variance. *Figure 5* reveals the well-known volatility clustering effect. In line with the developments in the CEE FX markets, some of which were mentioned in the previous section, the clusters are clearly evident during the first 250 days of the EUR/CZK and EUR/HUF sample periods, as they are during the last 120 days of the samples for all three exchange rates. The plots of autocorrelation functions in *Figure 4* (left column) corroborate this finding. In line with existing empirical evidence for other foreign exchange rates, the autocorrelation functions of the realized volatility for all three currencies decay very slowly, which is consistent with long-memory type of dynamics.

5.2 A Model for Daily Returns and Realized Variance

Motivated by the empirical results reported in the previous sections, we now turn to modeling the joint behavior of daily returns and realized variance. A successful empirical model must be able to capture the distributional and dynamic properties of returns and volatility observed in the data. At the same time, it should be sufficiently parsimonious to avoid issues associated with over-fitting and complicated estimation procedures.

After initial experimentation, we propose the following model:

Figure 5 Plots of Daily Realized Variance (*RV*), *left*, and of Daily Standardized Returns ($r_t/(RV)^{1/2}$), *right*
for the case of EUR/CZK (first row), EUR/HUF (second row), and EUR/PLN exchange rate returns (third row)

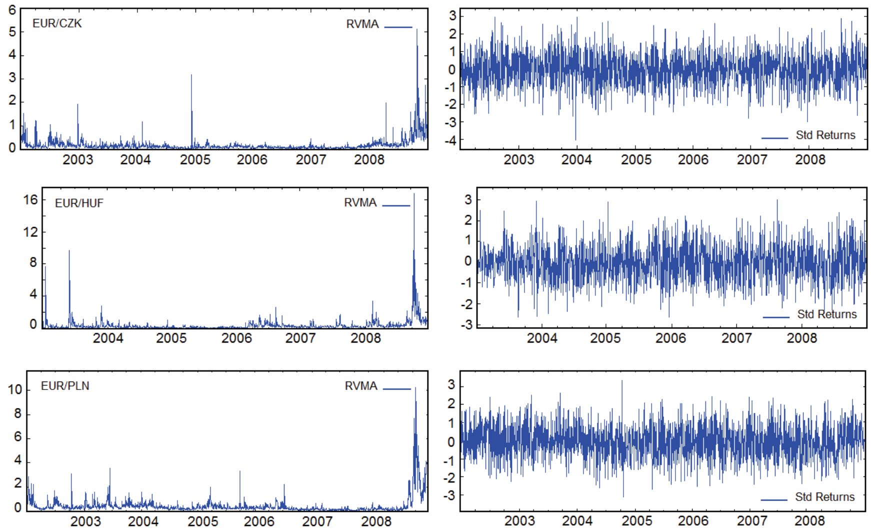

$$r_t = \mu + \sqrt{RV_t}\varepsilon_t \tag{11}$$

$$\log(RV_t) = \gamma_0 + \gamma_1 \log(RV_{t-1}) + \gamma_2 \log(RV_{t-1}^{(5)}) + \gamma_3 \log(RV_{t-1}^{(22)}) + \sqrt{h_t}\eta_t \tag{12}$$

$$h_t = \omega + \alpha\eta_{t-1}^2 + \beta h_{t-1} \tag{13}$$

where log $(RV_t^{(k)})$ denotes the average logarithmic realized variance over the past k days, i.e.

$$\log(RV_t^{(k)}) = \frac{1}{k}\sum_{j=1}^{k}\log(RV_{t-j+1})$$

and $\varepsilon_t \sim N(0,1)$, $\eta_t \sim \text{SkewedT}(\lambda, v)$, both *i.i.d.* The model has three equations. The first equation describes the evolution of daily returns. Since we found no serial correlation in the daily return series, we do not include any dynamics in the mean equation. Consistent with the observation that the daily returns standardized by daily realized volatility are approximately Gaussian, we assume that the return innovations follow the normal distribution.

The second equation represents the well-known heterogeneous autoregressive model (HAR) for logarithmic realized variance originally proposed by Corsi (2009). While not a genuine long-memory model, the HAR model captures remarkably well the persistent dynamics typically found in the time-series of realized variances across different asset classes (see Andersen, Bollerslev and Dobrev, 2007, Corsi, Mittnik, Pigorsch and Pigorsch, 2008, and Bollerslev, Kretschmer, Pigorsch and Tauchen, 2009, among others). Writing the model in terms of the logarithm of realized variance as opposed to realized variance itself has the obvious advantage of avoiding issues associated with possibly negative coefficient estimates.

Following Corsi, Mittnik, Pigorsch and Pigorsch, 2008, we generalize the model by allowing for conditional heteroskedasticity of the volatility innovations. This allows us to capture the volatility-of-volatility effect, i.e. the empirical observation that the volatility of volatility tends to increase when volatility itself increases. In the interest of parsimony, we adopt a simple GARCH(1,1) specification for the conditional variance of the logarithmic realized variance and let the innovation process follow the skewed Student-*t* distribution proposed by Hansen (1994). Again, a specification test will be provided to check the adequacy of this assumption.

We employ the method of maximum likelihood to estimate the parameters of the model. We obtain initial consistent estimates by estimating the mean equation separately from the equation for realized variance. This is equivalent to joint maximum likelihood estimation under the assumption of independence between ε_t and η_t.

Table 3 reports the maximum likelihood estimates for both the mean and the realized variance equations along with the corresponding standard errors and *p*-values as well as a set of specification tests for the residuals from the HAR-GARCH part of the model.

Starting with the HAR equation, we observe that the coefficient estimates on daily, weekly, and monthly variance components are all highly significant, a finding that corresponds to the results obtained in the previous applications of the HAR model for realized variance in the literature. In case of EUR/HUF, two lags of the logarithm of realized variance are used in equation (12) to improve the overall fit of the model. In terms of magnitude, the relative impact of the daily, weekly and monthly variance components differs across the exchange rates. Specifically, in case of EUR/CZK the monthly variance component seems to have the largest impact on current realized variance followed by the weekly and daily components. However, it is the daily component for EUR/HUF and the weekly component for EUR/PLN that seem to affect the current realized variance of the respective currencies the most.

The dynamics of the relative impact of different variance components on the current (realized) variance carries important information about the developments in the attitudes of the market participants towards short-, medium-, and long-term variance. For example, an upward trend present in the coefficient estimates on the long--term variance component informs us of an increasing degree of persistence of the long-term uncertainty in the market. In parallel, this may lead to a growing influence of the long-term term volatility on the short term volatility (Müller, Dacorogna, Dav, Olsen, Pictet and von Weizsacker, 1997). The economic intuition here is simple: as long-term volatility directly affects the expectations about the future market trends and risk, the short-term FX market participants use the information that it contains to adjust their trading behavior, thereby causing the volatility to increase in the short-term (Corsi, 2009).

To see whether the parameters of the HAR model are stable over time or if they undergo some structural changes, we plot in *Figure 6* recursive parameter estimates with 95% confidence intervals. For each of the three exchange rates we observe a clear tendency for the coefficient estimates on the monthly variance component to

Table 3 Estimation Results for the HAR-GARCH Model

Parameter estimates and diagnostics for the HAR-GARCH model. In panels (A)–(D), all coefficients are significant at 1% level, except where denoted by superscript *b* (significant at 5%), or *d* (insignificant at conventional levels). In panel (E), the diagnostic statistics pertain to the scaled residuals ($\hat{\varepsilon}_t, \hat{\eta}_t$) from the HAR-GARCH model.

$Q(20)$ and $Q^2(20)$ represent the Ljung-Box *Q*-statistics for the null hypothesis of no autocorrelation up to lag 20 in the raw and squared standardized residuals. Similarly, *LM* (20) represents Engle's LM test for ARCH effects up to lag 20 in the standardized residual series. R_j represents Patton's (2006) goodness-of-fit test statistic for correct specification of the dependence structure of ($\hat{\varepsilon}_t, \hat{\eta}_t$). The corresponding *p*-values are given in parentheses.

	CZK		HUF		PLN	
	Coeff.	S.E.	Coeff.	S.E.	Coeff.	S.E.
			A. Mean Equation			
μ	-0.010^d	(0.010)	0.008^d	(0.015)	0.009^d	(0.014)
			B. HAR Equation			
γ_0	-0.151	(0.045)	-0.108	(0.034)	-0.058^b	(0.025)
γ_1	0.243	(0.029)	0.371	(0.031)	0.234	(0.028)
γ_2	–		–		0.094	(0.029)
γ_3	0.298	(0.049)	0.294	(0.047)	0.381	(0.055)
γ_4	0.392	(0.045)	0.271	(0.040)	0.246	(0.037)
			C. GARCH Equation			
ω	0.070	(0.026)	0.062^b	(0.032)	0.014^d	(0.010)
α	0.091	(0.027)	0.070	(0.027)	0.029^b	(0.013)
β	0.668	(0.107)	0.749	(0.098)	0.910	(0.052)
λ	0.124	(0.036)	0.245	(0.041)	0.213	(0.036)
v	10.42	(2.087)	9.018	(2.199)	8.881	(1.798)
			D. Dependence Structure			
θ	–		0.278	(0.035)	0.124	(0.024)
			E. Diagnostics			
R^2	0.549		0.626		0.666	
$Q(20)$	26.20	(0.159)	21.20	(0.385)	16.64	(0.676)
$Q^2(20$	24.62	(0.136)	23.78	(0.162)	20.26	(0.318)
$LM(20)$	1.212	(0.234)	1.265	(0.193)	0.944	(0.530)
R_j	–		4.003	(0.779)	2.960	(0.889)

increase over time, while the other coefficients exhibit relatively stable behavior. This implies that the persistence of the realized variance increases over time: while towards the beginning of the sample the monthly component is small and statistically insignificant, it gradually increases and becomes highly significant as we add more recent observations. The dynamics of realized variance of the CE exchange rates thus increasingly resembles those of major exchange rates such as EUR/USD and USD/ /JPY.

Turning to the variance equation in the HAR-GARCH model, we note highly significant estimates of ARCH and GARCH coefficients for all three exchange rates that capture the clustering of volatility of realized volatility. In particular, the GARCH coefficient estimate informs us of a relatively large and positive effect of the previous period volatility on the current volatility of realized volatility. The highly significant

Figure 6 Recursive Parameter Estimates from the HAR Equations

Top, middle and bottom parts of the figures correspond to the EUR/CZK, EUR/HUF and EUR/ /PLN exchange rates, respectively. In each figure, the center line represents the parameter estimates, while the two outer lines depict the 95% confidence interval. We use log (*RV*)[-1] and log (*RV*)[-2] to denote the first and the second lag of the logarithmic realized variance, respectively and log (*RV*)[5] and log (*RV*)[22] the 5-day and the 22-day logarithmic realized variances. There are 1,508, 1,235, and 1,490 coefficient estimates for the EUR/CZK, EUR/HUF and EUR//PLN, respectively.

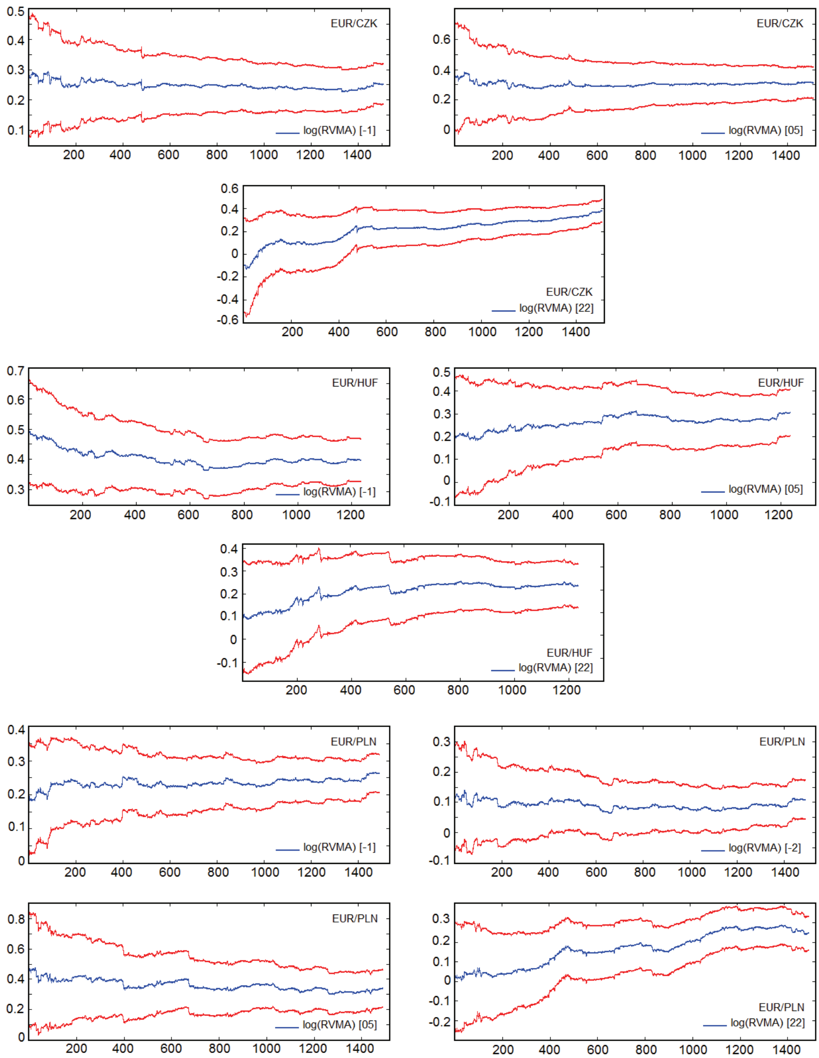

estimates of the asymmetry and the tail coefficients for the skewed Student-*t* distribution provide a preliminary indication of the validity of our initial assumption about the shape of the distribution of HAR innovations. In particular, the asymmetry parameter is positive implying positive skewness of the HAR innovations. The estimated numbers of degrees of freedom fall in the 8.9 to 10.4 range for the three exchange rates, implying that the HAR innovations exhibit substantially thicker tails than the normal distribution.

The residual diagnostics performed on the simple and squared standardized residuals from the HAR-GARCH equations confirm that our model provides an adequate fit to the data. Specifically, the Ljung-Box statistics verifies that neither raw nor squared residuals are serially correlated while Engle's LM test provides evidence of no remaining ARCH effects in the residual series.

We check the adequacy of the assumption of a skewed Student-t distribution for the innovation term in the HAR equation by plotting the kernel density estimate for the standardized residuals from the HAR-GARCH part of the model against a skewed Student-t density implied by the corresponding parameter estimates (*Figure 7*, left column). We find that for each of the three exchange rates, the latter provides a nearly perfect match to the residual kernel density estimate demonstrating that the assumption of a skewed Student-t distributed error term is indeed legitimate.

Finally, we examine the validity of the assumption that the innovation terms in the return and HAR equations (ε_t, η_t) are independent. Contrary to our initial assumption we find a small but statistically significant dependence between the two series in case of EUR/HUF and EUR/PLN. This means that the periods of depreciation of PLN and HUF w.r.t. EUR are associated with high *unexpected* volatility. In case of EUR/CZK, the linear correlation coefficient is statistically indistinguishable from zero.

To get an idea about the structure of dependence between the return and volatility innovations, we show in *Figure 7* (right column) a scatter plot of the residuals transformed into uniform variates by their respective estimated marginal distribution functions; that is, we plot $\hat{\mu}_{1,t} = \hat{F}_\varepsilon(\hat{\varepsilon}_t)$ against $\hat{\mu}_{2,t} = \hat{F}_\eta(\hat{\eta}_t)$. This transformation ensures that we focus on the dependence structure of (ε_t, η_t), free from the effects of the marginal distributions. If the two innovation processes (ε_t, η_t) are independent, ($\hat{\mu}_{1,t}, \hat{\mu}_{2,t}$) should be approximately uniformly distributed on $[0,1]^2$. The scatter plots show that this is the case of EUR/CZK. For EUR/HUF and EUR/PLN we observe slight positive dependence, which appears to be asymmetric in the former case: the innovations seem to be more dependent in the upper tail (upper right corner) than in the lower tail. Thus large depreciations of HUF tend to be accompanied by large unexpected volatility.

To incorporate these features into our model, we describe the dependence structure of the two innovation terms, ε_t and η_t, using copulas. By the theorem due to Sklar (1959), any bivariate distribution function $F_{x,y}(w_x, w_y)$ with continuous marginal distributions, $F_x(w_x)$, $F_y(w_y)$ can be written as

$$F_{x,y}(w_x, w_y) = C(F_x(w_x) + F_y(w_y))$$

where $C(\underline{u},v)$, $(u,v) \in [0,1]^2$, is a copula function satisfying certain properties (see e.g. Cherubini, Luciano and Vecchiato (2004) for an introduction to copulas in the context of financial modeling). Conversely, any copula together with arbitrary continuous marginal distribution functions yields a proper bivariate distribution

Figure 7

Left: Kernel density estimates for the standardized residuals from the HAR-GARCH equations vs. skewed Student-*t* density implied by the parameter estimates. *Right*: Scatter plots of standardized returns (*y*-axis) vs. standardized residuals from the HAR-GARCH equations, both transformed into uniform marginals by empirical CDF. The first (second, third) row corresponds to the case of EUR/CZK (EUR/HUF, EUR/PLN) exchange rate returns, respectively.

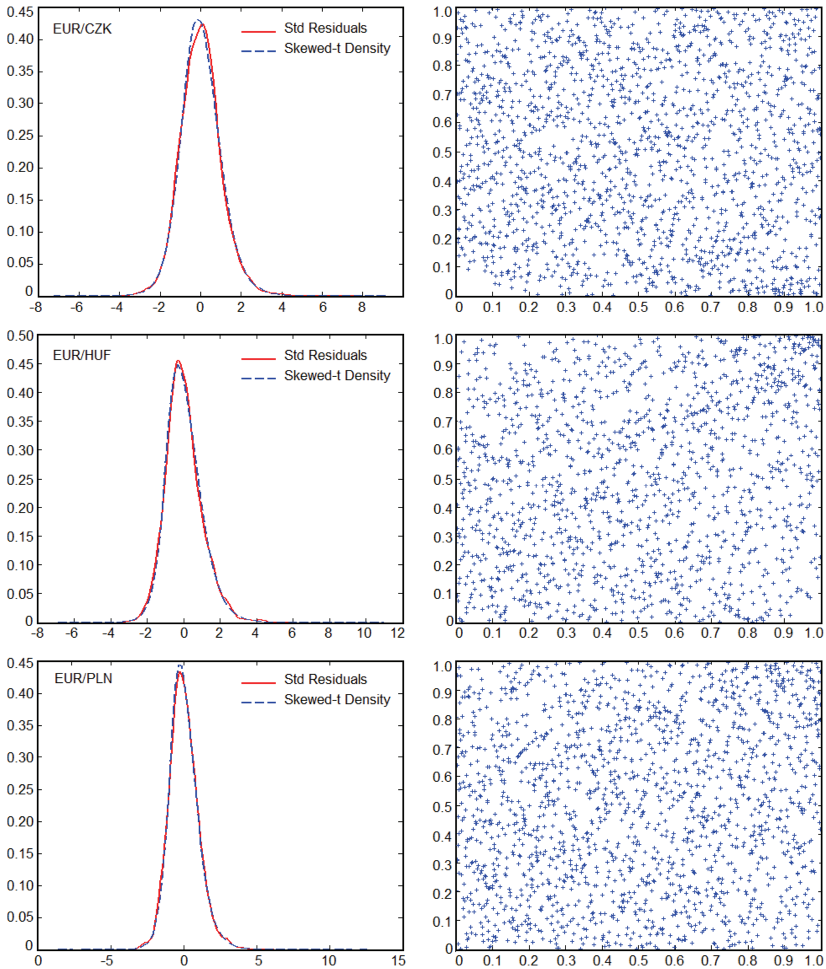

function. It is this latter property that we exploit here: given the excellent fit of the normal and skewed Student-*t* distributions for ε_t and η_t, respectively, we select an appropriate copula to link them together to obtain their joint distribution.

Following standard testing procedures, we find that the rotated Clayton and the Gaussian copulas, described in greater detail in the *Appendix*, provide suitable characterizations of the dependence structures of the EUR/HUF and EUR/PLN innovations, respectively. The maximum likelihood parameter estimates for the corresponding copula models are reported in Panel D of *Table 3*, along with a goodness-

-of-fit test due to Patton (2006) used to assess their statistical adequacy. We find that the dependence between the bivariate innovations is well described by the proposed copula specifications. The estimated parameters are statistically significant but indicate rather week dependence and hence the loss in efficiency from estimating our model equation-by-equation is probably quite small. Although straightforward, we do not further pursue the joint estimation of the model for this reason.

5.3 Forecasting Exercise

Given the satisfactory performance of the HAR-GARCH model in-sample, we proceed to evaluate its performance out-of-sample. As part of the analysis, we also investigate whether explicitly allowing for conditional heteroskedasticity in the innovations of realized volatility improves on the accuracy of the simple HAR model. This is motivated by the fact that despite yielding better fit in-sample, the HAR-GARCH model entails more parameters and hence potentially more serious parameter uncertainty problem than a simple HAR model, which may in turn adversely affect its forecasting performance.

We employ a Mincer-Zarnowitz (1969) regression (MZ) to assess the forecasting performance of the individual models. The MZ regression involves regressing the realized variance for time t, RV_t, on a constant and the volatility forecast at time t obtained using the information available at time $t-1$, $h_{t|t-1}$. Thus we estimate

$$RV_t = \alpha + \beta h_{t|t-1} + e_{t|t-1} \tag{14}$$

If the forecasting model performs well, the forecast is unbiased and the error of the forecast is small; in other words, $\alpha = 0$ and $\beta = 1$, and the R^2 implied by (16) is high. In our study we assess the performance of the models in forecasting logarithmic realized variance (obtained directly from equation (12)) as well as the squared root of realized variance and the realized variance itself, both obtained from the model for the log (RV_t) by taking the appropriate exponential transformation.

To help us differentiate between the forecasting performances of the HAR-GARCH vs. simple HAR models, we rely on two parametric loss functions, MSE and QLIKE, defined as:

$$MSE: L(RV_t, h_{t|t-1}) = (RV_t - h_{t|t-1})^2 \tag{15}$$

$$QLIKE: L(RV_t, h_{t|t-1}) = (\log(h_{t|t-1}) + RV_t) \tag{16}$$

Now commonly applied in the volatility forecasting literature, both MSE and QLIKE are known to deliver consistent rankings of realized variance forecasts when a noisy, but conditionally unbiased proxy is used in place of latent volatility (Patton, 2008). In addition, we note that while MSE penalizes both the positive and the negative forecast errors equally, the QLIKE imposes larger penalty when the volatility forecast underestimates the realized quantity, so that using the latter is of interest if underestimating future volatility is more costly.

Table 4 reports the coefficient estimates from the Mincer-Zarnowitz regressions based on the forecasts of log (RV), RV, and $(RV)^{1/2}$ obtained from HAR and HAR-GARCH models for logarithmic realized variance. As in case of the forecast evaluations further in the text, the regressions are based on 250 forecasts of daily

Table 4 Evaluation of HAR and HAR-GARCH Out-of-Sample Forecasts

Estimated intercept (α) and slope (β) coefficient in the Mincer-Zarnowitz regression of realized variance on its forecast, with heteroskedasticity consistent standard errors in parentheses. R^2 denotes the usual coefficient of determination in a linear regression. MSE and QLIKE are calculated according to formulas (15) and (16), respectively. Superscripts *b* (*c*) at the MSE and QLIKE denote the cases when the model produces significantly better forecasts than its competitor at 5% (10%) levels, respectively, based on the Giacomini-White test of equal predictive accuracy. The evaluation is based on 250 rolling forecasts of the logarithmic realized variance, $\log(RV_t)$, realized variance, RV_t, and realized volatility, $(RV_t)^{1/2}$, for each of the three exchange rates.

	CZK		HUF		PLN	
	HAR	HAR-G	HAR	HAR-G	HAR	HAR-G
			A. Forecasts of RV_t			
α	0.037	0.030	-0.028	-0.052	0.061	0.059
	(0.043)	(0.044)	(0.091)	(0.098)	(0.043)	(0.044)
β	1.223	1.254	1.400	1.454	1.209	1.219
	(0.192)	(0.198)	(0.238)	(0.254)	(0.128)	(0.134)
R^2	0.412	0.411	0.448	0.448	0.654	0.650
MSE	0.214	0.216	1.628	1.650	0.907	0.923
QLIKE	0.064	0.064	0.562	0.563	0.134^b	0.136
			B. Forecasts of $(RV_t)^{1/2}$			
α	0.030	0.021	-0.010	-0.023	-0.011	-0.012
	(0.047)	(0.048)	(0.065)	(0.068)	(0.032)	(0.032)
β	1.054	1.072	1.113	1.134	1.105	1.108
	(0.101)	(0.103)	(0.110)	(0.114)	(0.067)	(0.070)
R^2	0.555	0.555	0.619	0.618	0.808	0.805
MSE	0.045	0.046	0.105	0.106	0.076	0.077
QLIKE	0.537	0.538	0.763	0.765	0.550^b	0.552
			C. Forecasts of $\log(RV_t)$			
α	0.092	0.106	0.085	0.095	0.115	0.118
	(0.091)	(0.092)	(0.064)	(0.065)	(0.064)	(0.065)
β	0.966	0.975	0.974	0.982	1.048	1.046
	(0.057)	(0.058)	(0.054)	(0.056)	(0.035)	(0.036)
R^2	0.610	0.609	0.659	0.658	0.839	0.839
MSE	0.296	0.297	0.327	0.329	0.293^c	0.294

realized variance obtained for each of the three exchange rates via a rolling forecasting scheme. Specifically, we use the first $T = 1{,}508$ (EUR/CZK), 1,235 (EUR//HUF) and 1,490 (EUR/PLN) observations of daily realized variance to obtain the forecast for $T + 1$. The remaining 249 forecasts are then obtained by rolling the fixed estimation window forward and re-estimating the parameters of the model each time.

We observe that in case of both HAR and HAR-GARCH models, the forecasts produced by the models are unbiased across the three exchange rates, with the estimated α's being statistically indistinguishable from zero and the estimated β's being approximately equal to one. The only exception is the forecast of the logarithmic realized variance for EUR/PLN, but even there the intercept is significantly different from zero only marginally. We add that the same results hold for the three forms of realized variances being forecasted (Panels A to C).

We next notice a relatively high explanatory power (measured by R^2) across the models. Discussed in terms of the form of realized variance being forecasted, the models achieve the best results with the logarithmic realized variance (Panel A), in which case the R^2 is found to be just over 0.60 for EUR/CZK and nearly 0.84 for the EUR/PLN case. The explanatory power of the models seem to deteriorate by 5 and 30 percentage points on average in case of the forecasts of realized volatility (Panel C) and the realized variance (Panel B), respectively. The relatively worse performance for the realized variance forecasts is hardly surprising, given the fact that the time series of RV_t exhibits several "outliers" associated with periods of high volatility and/or potential jumps in the exchange rates. These outliers tend to be attenuated by taking the square root and especially logarithmic transformations resulting into better forecasting performance.

Comparing the relative performance of the simple HAR versus the more elaborate HAR-GARCH model, we observe that the former provides consistently albeit marginally better forecasting power across all three exchange rates and loss functions, the only exception being the forecasts of realized variance (Panel B) for EUR/HUF. To see whether the difference between the competing models is statistically significant we employ a test developed by Giacomini and White (2006). Note that this test is valid despite the fact that the two models are nested. This is due to the non-vanishing parameter estimation error implied by the rolling forecasting scheme, which prevents the test statistics from degenerating in the limit. Based on the Giacomini-White test we find that the HAR and HAR-GARCH perform equally well in forecasting the various volatilities of EUR/CZK and EUR/HUF. Some statistically significant difference is detected for EUR/PLN when the QLIKE loss function is employed, with the simple HAR beating the more complicated HAR-GARCH.

To summarize our forecasting exercise, we find that despite the HAR-GARCH model having a better in-sample fit, the simple HAR model offers equally or in some cases even significantly better forecasting performance. The fact that the simple HAR can be estimated by ordinary least squares makes it a particularly attractive forecasting model.

6. Conclusion

Our study extends the current understanding of the Central European exchange rates behavior by describing the conditional distribution and dynamics of EUR/CZK, EUR/HUF, and EUR/PLN exchange rate returns and volatility in the period from 2002 to 2008. Relying on model-free nonparametric measures of ex-post volatility based on the use of 5-minute intraday returns, our approach contrasts with the existing literature that almost invariably employs a parametric framework to model the exchange rate volatility.

Our findings show that daily returns on the EUR/CZK, EUR/HUF and EUR/ /PLN exchange rates are approximately normally distributed and independent over time, when properly scaled by model-free estimates of daily volatility. Given the properties of the 5-minute intraday returns, we find that a relatively simple correction to the realized variance suffices to account for the bias arising from the microstructure noise contaminating the data. The resulting daily realized variance exhibits substantial positive skewness and very persistent, long-memory type of dynamics.

We estimate a simple time series model for daily returns, realized variance and the time-varying volatility of realized variance. We show that the particular specification of the model that we suggest captures very well all salient features of the data and can be successfully employed for constructing point as well as density forecasts of future volatility. It can also serve very well as an auxiliary model for estimating stochastic volatility models often employed in derivatives pricing. The results from an out-of-sample forecasting exercise provide evidence of excellent forecasting performance of the HAR-GARCH model, especially in forecasting the logarithmic realized variance. It remains to be noted that a simple and computationally less demanding HAR model performs at least as well as and sometimes even better than the HAR-GARCH model.

Our findings provide a natural starting point for future investigation of the Central European exchange rates. The flexible and computationally simple non-parametric approach for measuring ex-post volatility that we employ can be used in areas ranging from volatility forecasting, to testing the efficiency of central bank's intervention along the lines of Beine, Lahaye, Laurent, Neely and Palm (2006), or analyzing the response of the volatility of the exchange rate to macroeconomic announcements (jumps). The simple, but highly empirically successful model for daily returns and volatility we propose in this paper could be employed to investigate and compare alternative continuous-time models and their ability to accurately price derivative securities written on the Czech koruna, Hungarian forint and Polish złoty exchange rates.

APPENDIX

Copulas

In Section 5.2 we employ the rotated Clayton and the Gaussian copulas to model the dependence between the two innovation processes in the HAR-GARCH model. The rotated Clayton copula $C_{RC}(u,v|\theta)$ is given by

$$C_{RC}(u,v|\theta) = u + v - 1 + (u^{-\theta} + v^{-\theta} - 1)^{-1/\theta}$$

where $\theta \in [0,\infty)$ governs the degree of dependence. The structure of dependence implied by the rotated Clayton copula is asymmetric in the sense that upper-tail extreme events are more dependent than lower-tail extremes.

The Gaussian copula $C_G(u,v|\theta)$ reads

$$C_G(u,v|\theta) = \Phi_\theta\left(\Phi^{-1}(u), \Phi^{-1}(v)\right)$$

where Φ_θ denotes the joint distribution function of a bivariate standard normal vector with correlation θ and Φ denotes the univariate standard normal distribution function. Contrary to the rotated Clayton copula, the dependence structure associated with the Gaussian copula is symmetric.

REFERENCES

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

Andersen TG, Bollerslev T, Diebold FX, Labys P (2000): Exchange Rate Returns Standardized by Realized Volatility Are (Nearly) Gaussian. *Multinational Finance Journal*, 4(3-4):159–179.

Andersen TG, Bollerslev T, Diebold FX, Labys P (2001). The Distribution of Realized Exchange Rate Volatility. *Journal of the American Statistical Association*, 96(453):42–55.

Andersen TG, Bollerslev T, Diebold FX, Labys P (2003): Modeling and Forecasting Realized Volatility. *Econometrica*, 71(2):579–625.

Andersen TG, Bollerslev T, 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.

Andersen TG, Bollerslev T, Frederiksen PH, Nielsen M (2009): Continuous Time Models, Realized Volatilities, and Testable Distributional Implications for Daily Stock Returns. J*ournal of Applied Econometrics* (forthcoming).

Antal J, Holub T (2007): Exchange Rate Arrangements prior to Euro Adoption. *Czech Economic Review*, 1(3):312–323.

Back K (1991): Asset Pricing for General Processes. *Journal of Mathematical Economics*, 20(4): 371–395.

Bandi F, Russell JR (2006): Separating Microstructure Noise from Volatility. *Journal of Financial Economics*, 79(3):655–692.

Barndorff-Nielsen OE, Hansen PR, Lunde A, Shephard N (2008): Designing Realized Kernels to Measure the Ex-Post Variation of Equity Prices in the Presence of Noise. *Econometrica*, 76(6): 1481–1536.

Barndorff-Nielsen OE, Shephard N (2002): Econometric Analysis of Realized Volatility and Its Use in Estimating Stochastic Volatility Models. *Journal of the Royal Statistical Society*, Series B, 64(2): 253–280.

Beine M, Lahaye J, Laurent S, Neely CJ, Palm FC (2006): Central Bank Intervention and Exchange Rate Volatility, its Continuous and Jump Components. *Federal Reserve Bank of St. Louis Working Paper*, 031B.

Bollerslev T, Kretschmer U, Pigorsch C, 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.

Bondarenko O (2007): *Variance Trading and Market Price of Variance Risk*. University of Illinois at Chicago.

Bontemps C, Meddahi N (2005): Testing Normality: A GMM Approach. *Journal of Econometrics*, 124(1):149–186.

Bulíř A, Šmídková K (2005): Exchange Rates in the New EU Accession Countries: What Have We Learned from the Forerunners. *Economic Systems*, 29(2):163–186.

Chernov M, Gallant AR, Ghysels E, Tauchen G (2003): Alternative Models for Stock Price Dynamics. *Journal of Econometrics*, 116(1-2):225–257.

Cherubini U, Luciano E, Vecchiato W (2004): *Copula Methods in Finance*. John Wiley & Sons.

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

Corsi F, Mittnik S, Pigorsch C, Pigorsch U (2008): The Volatility of Realized Volatility. *Econometric Reviews*, 27(1):46–78.

Égert B (2007): Central Bank Interventions, Communication and Interest Rate Policy in Emerging European Economies. *Journal of Comparative Economics*, 35(2):387–413.

Égert B, Komárek L (2006): Foreign Exchange Interventions and Interest Rate Policy in the Czech Republic: Hand in Glove? *Economic Systems*, 30(2):121–140.

Fidrmuc J, Horváth R (2008): Volatility of Exchange Rates in Selected New EU Members: Evidence from Daily Data. *Economic Systems*, 32(1):103–118.

Frömmel M (2007): Volatility Regimes in Central and Eastern European Countries' Exchange Rates. *Universiteit Gent Working Paper*, 2007/487.

Geršl A (2004): Foreign Exchange Intervention: The Theoretical Debate and the Czech Koruna Episode. *Finance a úvěr-Czech Journal of Economics and Finance*, 54(3-4):94–116.

Geršl A (2006): Testing the Effectiveness of the Czech National Bank's Foreign-Exchange Interventions. *Finance a úvěr-Czech Journal of Economics and Finance*, 56(9-10):398–415.

Geršl A, Holub T (2006): Foreign Exchange Interventions under Inflation Targeting: The Czech Experience. *Contemporary Economic Policy*, 24(4):475–491.

Giacomini R, White H (2006): Tests of Conditional Predictive Ability. *Econometrica*, 74(6):1545–1578.

Hansen BE (1994): Autoregressive Conditional Density Estimation. *International Economic Review*, 35(3):705–730.

Hansen PH, Lunde A (2006): Realized Variance and Market Microstructure Noise. *Journal of Business & Economic Statistics*, 24(2):127–161.

Hansen PR, Large J, Lunde A (2008): Moving Average-Based Estimators of Integrated Variance. *Econometric Reviews*, 27(1):79–111.

Hong Y, Li H (2005): Nonparametric Specification Testing for Continuous-time Models with Application to Spot Interest Rates. The *Review of Financial Studies*, 18(1):37–84.

Horváth R (2005): Exchange Rate Variability, Pressures and Optimum Currency Area Criteria: Lessons for the Central and Eastern European Countries. *Czech National Bank Working Paper*, no. 8/2005.

Kočenda E, Valachy J (2006): Exchange Rate Volatility and Regime Change: A Visegrád Comparison. *Journal of Comparative Economics*, 34:727–753.

McAleer M, Medeiros MC (2008): Realized Volatility: A Review. *Econometric Reviews*, 27(1-3): 10–45.

Melecký M, Komárek L (2008): Transitional Appreciation of Equilibrium Exchange Rates and the ERM II. *Transition Studies Review*, 15(1):95–110.

Müller UH, Dacorogna MM, Dav RD, Olsen RB, Pictet O, Weizsacker JE von (1997): Volatilities of Different Time Resolutions – Analysing the Dynamics of Market Components. *Journal of Empirical Finance*, 4(2):213–239.

O'Hara M (1995): *Market Microstructure Theory*. Blackwell Publishing.

Oomen RCA (2006): Properties of Realized Variance under Alternative Sampling Schemes. *Journal of Business & Economic Statistics*, 24(2):219–237.

Papaioannou MG (2006): Exchange Rate Risk Measurement and Management: Issues and Approaches for Firms. *International Monetary Fund Working Paper*, no. 06/255.

Patton AJ (2006): Modeling Asymmetric Exchange Rate Dependence. *International Economic Review*, 47(2):527–556.

Patton AJ (2008): *Volatility Forecast Comparison Using Imperfect Volatility*. Oxford-Man Institute of Quantitative Finance, University of Oxford.

Protter PE (2005): *Stochastic Integration and Differential Equations*. 2ed. Springer.

Sklar A (1959): Fonctions de Répartition à n Dimensions et Leurs Marges. *Publications de l'Institut de Statistique de l'Université de Paris*, 8:229–231.

Zhang L, Mykland PA, Aït-Sahalia Y (2005). A Tale of Two-Time Scales: Determining Integrated Volatility with Noisy High-Frequency Data. *Journal of the American Statistical Association*, 100(472):1394–1411.

Žikeš F (2008): *Testing for Normality and Independence of Asset Returns Standardized by Realized Measures of Integrated Variance*. Imperial College London.