UNIVERSITY OF GOTHENBURG
SCHOOL OF BUSINESS, ECONOMICS AND LAW

Master Degree Project in Finance

Forecasting Exchange Rate Volatility

Applying HAR models and Implied Volatility in SEK denominated markets

Anton Agermark and Visar Hoti

Supervisor: Adam Farago
Master Degree Project No. 2016:118
Graduate School

Abstract

In this paper we study a set of models' forecasting accuracy of realized volatility in two SEK denominated exchange rates, EUR/SEK and USD/SEK, with the purpose to analyze if ex-post or ex-ante forecasting models produce the most accurate forecasts. High-frequency exchange rate data is employed in order to construct the ex-post Heterogeneous Autoregressive Model of Realized Volatility, HAR-RV, as well as a modified model using the bipower and tripower variation to separate the continuous sample path (C) and the jump component (J) of realized volatility, HAR-CJ. The forecasting accuracy of the ex-ante implied volatility estimate (IV) is also evaluated, based on daily OTC data and regarded as the option market's forecast of future volatility. The forecasts are conducted applying in-sample and out-of-sample tests over two horizons, one week and one month. Our findings do not provide clear evidence whether to rely solely on the ex-ante or ex-post estimate when forecasting exchange rate volatility. Rather, the model combining ex-post and ex-ante information, HAR-RV-IV, consistently provides good forecasting results.

Keywords: *forecasting, implied volatility, realized volatility, jump process, bipower variation, tripower variation, high-frequency data, FX*

Acknowledgments

We wish to express our deepest appreciation to our supervisor Adam Farago, for his expertise and very generous guidance along the road.

Contents

List of Tables

List of Figures

1 Introduction

One of the most important determinants of risky assets in financial markets, volatility, has been extensively researched both empirically and theoretically. Ultimately, derivative and asset pricing, hedging, and risk management involves a valuation procedure assessing an asset's level of riskiness with reference to future payoffs. The ability to properly forecast future volatility with current information is therefore of particular importance. However, opinions differ in regard to what model produces the most accurate volatility forecasts where a fundamental issue concerns whether one should rely on ex-post or ex-ante measures when conducting forecasts.

Numerous studies examining ex-post volatility models find evidence of high predictability using different types of the Heterogeneous Autoregressive (HAR) models across different financial markets (e.g. Andersen et al. [2007], Corsi [2009], and Liu et al. [2015]). Another widely examined volatility forecasting procedure is to adopt the ex-ante implied volatility, representing the markets view of future volatility implied in option prices, as a predictor of future volatility. Studies using this measure find mixed results on its performance vis-à-vis the ex-post volatility estimates (see Jorion [1995], Christensen and Prabhala [1998], Martens and Zein [2004], and Busch et al. [2011]). The vast majority of studies on the topic of implied volatility have made use of exchange traded options from which implied volatility has been backed out using option pricing models, the most famous being the Black-Scholes-Merton option pricing formula. However, recent findings suggest using implied volatilities based on over-the-counter (OTC) at-the-money (ATM) options rather than exchange traded ones (see Li [2002], Kellard et al. [2010], and Pong et al. [2004]). The advantages of incorporating OTC options compared to exchange traded options are that the trade volumes of the former by far exceed the volumes of the latter and that ATM options ensure moneyness, thus reducing the risk of measurement errors. Moreover, exchange traded options only trade for seven major foreign currencies, all denominated in U.S. dollars (USD), in contrast to OTC options that have a market for hundreds of currency pairs. OTC options therefore enable us to study exchange rates denominated in the Swedish krona (SEK).

The purpose of this paper is to compare ex-post and ex-ante models' forecasting accuracy of realized volatility constructed using high-frequency data on two SEK denominated currency pairs. The procedure involves the application of in-sample and out-of-sample forecasts across two horizons, one week and one month. The evaluation is done assessing a goodness-of-fit measure for the in-sample forecasts, and analyzing two forecasting error statistics for the corresponding out-of-sample forecasts.

The motivation of this paper is essentially two-fold. First, to our best knowledge our study is the first to conduct volatility forecasting comparisons incorporating OTC implied volatility data in combination with HAR models calculated from high-frequency data. Second, previous studies focus exclusively on more actively traded exchange rates in the FX market which makes our study one of the first to conduct comparative volatility forecasting on SEK denominated exchange rates.

We sample high-frequency data, on five-minute intervals, over a period of more than eight years to construct a measure of realized volatility for the EUR/SEK and USD/SEK exchange rates. Utilizing this, the ex-post forecasts are executed using two types of Heterogeneous Autoregressive (HAR). We employ the simple Heterogeneous Autoregressive model of Realized Volatility, HAR-RV, proposed by Corsi (2009), and a modified version accounting for the continuous sample path (C) and the jump component (J) of realized volatility, HAR-CJ, proposed in Andersen et al. (2007). The latter is based directly on the theoretical results found in Barndorff-Nielsen and Shephard (2004) who introduced a procedure using bi- and tripower variation to separate C and J. Apart from evaluating whether to rely on ex-post or ex-ante measures when forecasting realized volatility, we further analyze if the ex-post forecasting models accuracy improves when modeling for the continuous sample path and the jump components. Furthermore, the ex-ante forecasts are done applying daily OTC implied volatility data based on options with a fixed maturity of one month. Lastly, two combined models incorporating implied volatility with the HAR models are employed in order to evaluate if a combination of ex-post and ex-ante measures strengthens the forecasting accuracy as well as to analyze if any of the measures are informationally efficient over the other.

Our main results cannot confirm whether ex-post or ex-ante forecasting models produce the most accurate forecast since the results deviate between the currency pairs and seem to depend on the employed methodological approach. However, across all tests the combined HAR-RV-IV model performs relatively well, indicating that a combination of ex-post and ex-ante information seems to create a favorable model.

1.1 Previous research

Whether implied or time-series volatility models produce the most accurate volatility forecasts has been subjected in a vast number of research papers (see Poon and Granger [2003] for a comprehensive review of the voluminous literature on volatility forecasting in financial markets). The wide literature has predominantly studied the issue concerning forecasts' based on historical realized volatility and that they should have no incremental explanatory power regarding the underlying asset's future volatility, if the options market is efficient and if the correct option pricing model is used. This model might for instance be the Black, Scholes and Merton option pricing model (BSM model) in Black and Scholes (1973), and Merton (1973). Turning forecasts into profitable trading strategies is otherwise possible if time-series volatility models induce auxiliary information for forecasting future volatility (Jorion [1995]; Hull [2014], p.435).

The time-series volatility forecasting research is to a large degree focused on the performance of ARCH(q) models, originally introduced in Engle (1982), only to continue with the more sophisticated and popular GARCH(1,1) model proposed by Bollerslev (1986). In an early study, Scott and Tucker (1989) use daily data ranging from March 1983 to March 1987 for exchange traded American currency call options on the British pound, Canadian dollar, Deutschemark, Japanese yen, and the Swiss franc against U.S. dollar to measure how well implied volatility perform over different term structures. Their findings display a coefficient of determination, which captures the informational content, in the region of 40-50 percent for six- to nine-month horizons with the conclusion that subsequent realized volatility is well captured by implied volatility and that no improvement on predictability is eminent when adding historical volatility to the regression.

Covering a period from 1985 through 1992 on three major currency pairs (franc, yen, and deutschemark against the dollar), Jorion (1995) compares how implied volatility performs relative to statistical time-series models in terms of informational content and predictive ability. Jorion finds that the implied volatility outperforms historical time-series volatility for all three currency pairs. Similarly, but with a slightly shorter time-series data set ranging from 1985 through 1991, Xu and Taylor (1995) examine whether exchange traded currency options on the pound, franc, yen, and deutschemark against the dollar are informationally efficient, and find that implied volatility contains incremental information relative to time-series forecasts about future exchange rate volatility. Fleming (1998) confirms Jorion's results that the implied volatility contains relevant information regarding future volatility incremental to time-series forecasts. Fleming argues in similar fashion as Christensen and Prabhala (1998) that incremental information found in historical time-series probably suffer from statistical artifacts caused by overlapping data usage.

The common aspect surrounding these studies is that all adopt low sampling frequencies such as daily returns to calculate realized volatility. More recent studies however identify the increased availability of high-frequency intraday data in time-series forecasts and argue that the predictability should improve by employing this data. Andersen and Bollerslev (1998) introduce the model-free realized volatility measure, defined as the sum of the squared intraday returns, and show that these produce the best realized volatility forecasts on the foreign exchange market. The volatility forecasting abilities on the pound, deutschemark and yen against the dollar are estimated by Li (2002) using high-frequency time-series data. Li finds that the latter provide incremental information to that of implied volatilities for the six-month forecast horizons. It is not as straightforward to infer the same about stock or stock index returns as shown by Blair et al. (2001) where they find that no significant incremental forecasting information can be distinguished by extending GARCH with high-frequency data. Martens and Zein (2004) note however that Li (2002) uses overlapping data, as did Canina and Figlewski (1993) and Lamoureux and Lastrapes (1993), which tend to favor time-series forecasts. Nevertheless, the outcome of the contest between implied volatilities and time-series models changes when the long memory characteristics from volatilities are implemented in forecasts built from squared high-frequency returns. What Martens and Zein (2004) find is that implied volatilities can be outperformed even using non-overlapping data, especially with regard to forecasting volatilities of the S&P 500 and Sweet Crude Oil, but the implied still performs better forecasts for the currency pair YEN/USD.

A comprehensive study by Pong et al. (2004) compare four methods to forecast realized volatility for the pound, deutschemark and yen against the dollar using OTC data and intraday returns with different horizons ranging from one day to three months. In contrast to Martens and Zein (2004) and Li (2002), who are unable to distinguish if the incremental information of historical forecasts emerge from the practice of a long memory model or from high-frequency returns, Pong et al. (2004) are able to find evidence that the usage of high-frequency data enhance the historical time-series forecasts rather than the selection of a long memory model by noting that the short memory and long memory forecasts yield similar outcomes. Moreover, their findings show that for the one-day and one-week horizons the intraday measures provide the most accurate forecasts, whereas implied volatilities are incremental to historical time-series forecast over the remaining long horizons. Thus, time-series forecasts prevail over the short horizon while implied volatility estimates dominate over long horizons. Applying the same method with over-the-counter option prices on currencies instead of exchange traded options, Christoffersen and Mazzotta (2005) also find evidence that implied volatility subsumes the information content from historical time-series and produces accurate one-month and three-month realized volatility forecasts.

A recent stage in the development of volatility forecasting is taken into account in Barndorff-Nielsen and Shephard (2004), Andersen et al. (2007), Corsi (2009), Busch et al. (2011), and Liu et al. (2015). They estimate volatility using the Heterogeneous Autoregressive model of Realized Volatility (HAR-RV) proposed by Corsi (2009) and inspired by the Heterogeneous Market Hypothesis from Müller et al. (1993). The HAR model recognizes traders' heterogeneous perceptions across markets and employs high-frequency intraday data by considering the linear cascade of moving averages derived over different time horizons. Corsi (2009) finds that the HAR-RV model outperforms more complex long-memory volatility models, however, using overlapping data in combination with HAR models cause severe correlation in the error terms according to Andersen et al. (2007). A modified version accounting for the continuous sample path (C) and the jump (J) component of realized volatility, HAR-CJ, was proposed in Andersen et al. (2007) based directly on the theoretical results found in Barndorff-Nielsen and Shephard (2004) that introduced a procedure using the bipower variation to separate C and J from the realized volatility. Comparing the implied volatility measure, backed out from option prices, against the HAR-RV and HAR-CJ models, Busch et al. (2011) is able to show that implied volatility contains incremental information about future volatility in the foreign exchange market. Finally, the HAR-RV model proves to be successful in the out-of-sample setting that Liu et al. (2015) expose the model for, without subjecting it to a comparison with implied volatility.

To summarize, early studies compared ex-ante implied volatilities backed out from exchange traded options against ex-post measures sampled at daily frequency, resulting in rather ambiguous conclusions. Academia shifted focus towards the comparison between daily data and high-frequency data as the latter became available, and the opportunity to build heterogeneous autoregressive models evolved. The overall picture concerning the adoption of overlapping vs. non-overlapping series favor employing the latter, and in order to avoid problems caused by overlapping data we estimate the models as functions of the non-overlapping parameters.

The remainder of the paper is organized as follows. Section 2 acquaints the reader with implied volatility and realized volatility, and describes the separation of the latter into its continuous sample path and jump components. Section 3 introduces the notation of the models along with their properties, and presents the two diverse forecasting methods deployed. Section 4 describes the employed data set, and Section 5 displays the empirical results. Section 6 concludes.

2 Theoretical framework

In this section we outline the theoretical framework for constructing the ex-ante implied volatility and the ex-post realized volatility measures. Initially, in Subsection 2.1, we present the implied volatility measure. Subsequently, in Subsection 2.2, we present the realized volatility measure. Finally, in Subsection 2.3, we provide a detailed explanation of the realized volatility separation into its continuous and jump components. For a even more detailed explanation regarding the realized volatility measure and its separation into the continuous and jump components, see e.g. Barndorff-Nielsen and Shephard (2004), Andersen et al. (2007), and Busch et al. (2011).

2.1 Implied volatility

Volatility estimates that are implied by option prices observed in the market are known as implied volatilities. The volatility is the sole parameter that cannot be directly observed in the Black-Scholes-Merton pricing formula (Hull [2014], p.318f) defined as:

$$c = S_0 N(d_1) - K e^{-rT} N(d_2), \tag{1}$$

$$p = K e^{-rT} N(-d_2) - S_0 N(-d_1), \tag{2}$$

$$d_1 = \frac{ln(\frac{S_0}{K}) + (r + \frac{\sigma^2}{2})T}{\sigma\sqrt{T}}, \tag{3}$$

$$d_2 = d_1 - \sigma\sqrt{T}. \tag{4}$$

The function *N(x)* is the cdf for the standard normal distribution. The values c and p are the European call and put prices, S_0 is the particular security's price at time zero, K is the strike price, r is the continuously compounded risk-free rate, T is the time to maturity of the option, and the one parameter that is not observable is σ, the volatility of the underlying security. Therefore, the value of σ that gives the price c or p is the implied volatility. An iterative process can be used to find the value of σ since there is no closed form solution to solve for it, i.e. it is not possible to invert equation (1) and equation (2) and express σ as a function of the observable parameters. Implied volatility is an ex-ante measure, i.e. forward looking, in contrast to historical volatility (ex-post) and is therefore used to monitor the market's expectation about a certain asset's future volatility. As other model-free methods exist to generate the implied volatility, it does not necessarily need to be backed out of the

Black-Scholes-Merton model (Hull [2014], p.418, and Li [2002]). Therefore, as described in Section 4 below, our ex-ante volatility measure does not imply the adoption of the Black-Scholes-Merton model.

2.2 Realized volatility

Corsi (2009) explains why realized volatility, if sampled correctly, is a consistent ex-post estimate of actual, often called integrated volatility. The model

$$dp(\tau) = \mu(\tau)dt + \sigma(\tau)dW(\tau) \tag{5}$$

represents a stochastic volatility process where $p(\tau)$ is the logarithm of the price, $\mu(\tau)$ is a continuous, finite variation process, $W(\tau)$ is a standard Brownian motion, and $\sigma(\tau)$ is a stochastic process independent of $W(\tau)$. For this process, the integrated variance on day t is the integral of the instantaneous variance over the one day interval

$$\sigma_t^2 = \int_{t-1}^{t} \sigma^2(w)dw. \tag{6}$$

Corsi highlights the findings in Andersen, Bollerslev, Diebold, and Ebens (2001), Andersen, Bollerslev, Diebold, and Labys (2001), and Barndorff-Nielsen and Shephard (2002a, 2002b), which shows that as sampling frequency increases towards the limit, using the discretely sampled and equally spaced intraday squared returns, the approximation of the true integrated variance becomes arbitrarily precise. This means that the integrated volatility of the Brownian motion can be approximated by the sum of the intraday squared returns. Busch et al. (2011) define this measure, realized variance, as

$$RV_t = \sum_{j=1}^{M} r_{t,j}^2 \qquad \text{for j= 1,..., M and t= 1,..., T,} \tag{7}$$

where $r_{t,j} = p_{t,j} - p_{t,j-1}$ are the intraday returns for day t with a sample frequency of j. Hence, realized variance is defined as the sum of squared intraday returns, and realized volatility is the square root of the realized variance, $(RV_t)^{1/2}$.

2.3 Modeling for the continuous sample path and jump component

Further assumptions will be made to account for jumps in the exchange rates. We apply the methodology of Andersen et al. (2007) and Busch et al. (2011)

when modeling the continuous and jump components. Equation (5) is now extended to the general stochastic volatility jump model taking the form

$$dp(\tau) = \mu(\tau)dt + \sigma(\tau)dW(\tau) + k(\tau)dq(\tau), \tag{8}$$

where the additional process $k(\tau)dq(\tau)$ that differentiates equation (8) from equation (5) is constructed by a counting process $q(\tau)$ that is normalized in the sense that $q(\tau) = 1$ in the presence of a jump at time τ and zero otherwise. Accordingly, $k(\tau)$ is the size of the jump at time τ given that $q(\tau) = 1$. A theoretical description concerning the separation of the continuous sample path and jump component of realized variance now follows. For the aforementioned process, the integrated variance for day t is the integral of the instantaneous variance plus the sum of squared jumps throughout the day:

$$\sigma_t^2 = \int_{t-1}^{t} \sigma^2(w)dw + \sum_{q(\tau)=1} k^2(\tau). \tag{9}$$

Hence, in the absence of jumps the integrated variance will be equal to instantaneous variance. In order to separate the continuous sample path and the jump component in equation (9) we make use of the related bipower and tripower variation measures. The realized bipower measure is

$$BV_t = \mu_1^{-2} \frac{M}{M-(k+1)} \sum_{j=k+2}^{M} |r_{t,j}||r_{t,j-k-1}|, \qquad j = 1, ..., M, \tag{10}$$

where $\mu_1 = \sqrt{2/\pi}$ is a constant, M is the total number of squared return observations at day t. According to Busch et al. (2011) who refers to Barndorff-Nielsen and Shephard (2007), and Hansen and Lunde (2006), a higher value of M improves precision of the estimators, but in practice it also makes the estimates more susceptible to market microstructure effects, such as bid-ask bounces, stale prices and measurement errors. The aforementioned studies show that setting k=1 decreases this bias. Next, the realized tripower quarticity is

$$TQ_t = \mu_{4/3}^{-3} \frac{M^2}{M-2(k+1)} \sum_{j=2k+3}^{M} |r_{t,j}|^{(4/3)} |r_{t,j-k-1}|^{(4/3)} |r_{t,j-2k-2}|^{(4/3)}, \tag{11}$$

where $\mu_{4/3} = 2^{2/3}\Gamma(7/6)/\Gamma(1/2)$ is based on the gamma function. Barndorff-Nielsen and Shephard (2004) and Barndorff-Nielsen, Shephard and Winkel

(2006) show that BV_t is a consistent estimator of the integrated variance part of equation (9) when M $\to \infty$ and accordingly, the difference $RV_t - BV_t$ converges towards the sum of squared jumps. Furthermore, Busch et al. (2011) make use of the ratio test statistic to distinguish what movement in the exchange rate during day t should be defined as a jump. Let us define the ratio test, specified as

$$Z_t = \sqrt{M}\frac{(RV_t - BV_t)RV_t^{-1}}{\left((\mu_1^4 + 2\mu_1^{-2} - 5)\max\left\{1, TQ_t BV_t^{-2}\right\}\right)^{1/2}}. \tag{12}$$

In the absence of jumps, $Z_t \to$ N(0,1) when M $\to \infty$. Large positive values of Z_t indicate that a jump has taken place and the limit is set in the jump equation below

$$J_t = I_{Z_t > \phi_{1-\alpha}}\left(RV_t - BV_t\right), \tag{13}$$

where I_X is the indicator for the event X. In order to detect jumps and to construct the series for J and C, we set the significance level $\alpha = 0.1\%$ as suggested by Andersen et al. (2007). If the Z-statistic exceeds $\phi_{1-\alpha}$ it indicates a jump and $I_{Z_t > \phi_{1-\alpha}} = 1$, which further translates to J_t being excess realized variance above bipower variation from equation (13). The outcome is then that a jump in the exchange rate is noted. The continuous component is defined as the remaining part of the quadratic variation:

$$C_t = RV_t - J_t. \tag{14}$$

Note from equation (14) that in the absence of a jump at day t the continuous part will be equal to the realized variance.

3 Methodology

In this section we aim to describe the econometric properties of the different applied models. We start of by considering the two ex-post volatility forecasting models: Subsection 3.1 displays the Heterogeneous Autoregressive model of Realized Volatility (HAR-RV), and Subsection 3.2 explains the HAR-CJ model separating the continuous (C) sample path and jump (J) components of realized volatility. Subsequently, Subsection 3.3 deals with the ex-ante forecasting measure of future realized volatility; implied volatility (IV). Further, the combined models HAR-RV-IV and HAR-CJ-IV are presented in Subsection 3.4. We then discuss the construction of the variables along with the manner in which we conduct the robustness test in Subsection 3.5. We conclude this section with Subsection 3.6 by presenting the two employed forecasting methods; in-sample and out-of-sample forecasts.

3.1 HAR-RV

Corsi (2009) proposes the HAR-RV model and shows that it, despite its simplicity, outperforms more complex models when forecasting volatility across different markets, including foreign exchange. Corsi explains that the model can be seen as a three-factor stochastic volatility model, where the factors are the past realized volatilities viewed at different frequencies. He motivates the use of three time frames by referring to the Heterogeneous Market Hypothesis (Müller et al. [1993]) that recognizes traders' heterogeneous perceptions across markets, i.e. short-term, medium-term and long-term investment horizons.

$$RV_{t+1,t+h} = \beta_0 + \beta_d RV_t + \beta_w RV_{t-4,t} + \beta_m RV_{t-21,t} + \epsilon_{t+1,t+h}, \qquad (15)$$

where RV_t is day t's realized volatility, defined as the square root of day t's realized variance, $RV_{t-4,t} = \left(\frac{1}{5}\sum_{k=0}^{4} RV_{t-k}\right)^{1/2}$ is the average of the previous trading week's realized volatility, $RV_{t-21,t} = \left(\frac{1}{22}\sum_{k=0}^{21} RV_{t-k}\right)^{1/2}$ is the average of the previous trading month's realized volatility, and $\epsilon_{t+1,t+h}$ is the forecasting error. Furthermore, when forecasting weekly and monthly volatility the left hand side of equation (15) will correspond to the average of the realized volatility for the coming five trading days, $RV_{t+1,t+5} = \left(\frac{1}{5}\sum_{k=1}^{5} RV_{t+k}\right)^{1/2}$ and twenty-two trading days, $RV_{t+1,t+22} = \left(\frac{1}{22}\sum_{k=1}^{22} RV_{t+k}\right)^{1/2}$ respectively.

3.2 HAR-CJ

Andersen et al. (2007) highlight that many log price processes are best described by a smooth and very slow mean-reverting continuous sample path process and a much less persistent jump component but that previous research concerning realized volatility forecasting has paid this relatively little attention. Following Andersen et al. (2007) we separate the RV regressors into their continuous (C) and jump (J) components and denote the model HAR-CJ. The notation is consistent with the specifications in the HAR-RV model in equation (15) with the slight difference that the disintegrated components J and C are regressed instead of past RV to forecast future RV.

$$RV_{t+1,t+h} = \beta_0 + \beta_d J_t + \beta_w J_{t-4,t} + \beta_m J_{t-21,t} + \beta_d C_t + \beta_w C_{t-4,t} + \beta_m C_{t-21,t} + \epsilon_{t+1,t+h}, \tag{16}$$

where J_t and C_t are the square root of the previous day's jump and continuous components, respectively. $J_{t-4,t} = \left(\frac{1}{5}\sum_{k=0}^{4} J_{t-k}\right)^{1/2}$ and $C_{t-4,t} = \left(\frac{1}{5}\sum_{k=0}^{4} C_{t-k}\right)^{1/2}$ are the average of the previous trading week's components. Moreover, $J_{t-21,t} = \left(\frac{1}{22}\sum_{k=0}^{21} J_{t-k}\right)^{1/2}$ and $C_{t-21,t} = \left(\frac{1}{22}\sum_{k=0}^{21} C_{t-k}\right)^{1/2}$ are the average of the previous trading month's components, and $\epsilon_{t+1,t+h}$ is the forecasting error. The left hand side of equation (16) is defined as in equation (15) above.

3.3 IV

The informational content of implied volatility and the evaluation of the volatility forecasting performance is carried out by applying the procedure of previous studies on the subject (see Christensen and Prabhala [1998], and Poon and Granger [2003]). The regression model takes the form,

$$RV_{t+1,t+h} = \beta_0 + \beta_{IV} IV_t + \epsilon_{t+1,t+h}, \tag{17}$$

where IV_t denotes the implied volatility measure at day t, and the left hand side variable $RV_{t+1,t+h}$ is defined as in equation (15).

3.4 Combined models: HAR-RV-IV & HAR-CJ-IV

As proposed by Busch et al. (2011), we modify the HAR-RV and HAR-CJ models by including IV and abbreviating them HAR-RV-IV and HAR-CJ-IV.

The HAR-RV-IV model is

$$RV_{t+1,t+h} = \beta_0 + \beta_d RV_t + \beta_w RV_{t-4,t} + \beta_m RV_{t-21,t} + \beta_{IV} IV_t + \epsilon_{t+1,t+h}, \quad (18)$$

where the variables are defined as previously described in equations (15) and (17). The HAR-CJ-IV model then follows as

$$\begin{aligned} RV_{t+1,t+h} =& \beta_0 + \beta_d J_t + \beta_w J_{t-4,t} + \beta_m J_{t-21,t} \\ &+ \beta_d C_t + \beta_w C_{t-4,t} + \beta_m C_{t-21,t} + \beta_{IV} IV_t + \epsilon_{t+1,t+h}, \end{aligned} \quad (19)$$

and the variables follow the same structure as in equations (16) and (17).

3.5 Variable construction and robustness tests

In Section 1.1 we highlighted the prevailing opinion in academia, following Christensen and Prabhala's (1998) findings, that it is preferable to use non-overlapping data when dealing with volatility time-series, as overlapping data has been shown to cause correlation in the error term when performing regressions. Consequently, we conduct the empirical tests using non-overlapping data and the manner in which it is constructed needs further explanation. The implied volatility measures have a constant time to expiration of one month and in order to create the non-overlapping time series, the ending and starting point of each forecasting period departure from the implied volatilities expiration dates. Building on this, the past realized volatility measure ends on the day of each expiration date and future realized volatility measure starts one day ahead of the expiration, i.e. the day after the new selected implied volatility one-month time to maturity starts. Figure 1 illustrates the procedure.

Figure 1: Description of variable construction

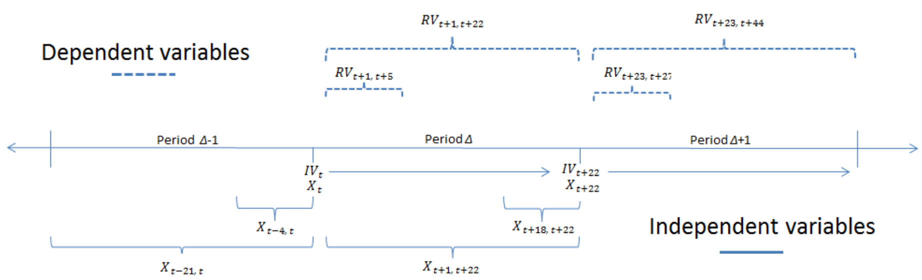

The dependent variables for the two forecasting horizons, one week and one month, are disclosed on the upper half where the forecasting period Δ is represented by $RV_{t+1,t+5}$ and $RV_{t+1,t+22}$. The independent variables used to forecast period Δ are the RV observations across three frequencies in period $\Delta - 1$ when applying the HAR-RV model and C and J observations when using the HAR-CJ model. The notation X is used for generalizing purposes, where $X = \{RV, J, C\}$. The IV observation is sampled at time t, forecasting $RV_{t+1,t+5}$ and $RV_{t+1,t+22}$ in period Δ.

In Section 5.3 a number of tests are conducted in order to verify the robustness of the original results. The first set of robustness tests addresses whether the form that the variables are specified in influences the results. Andersen et al. (2007) evaluates this by using realized variance, realized volatility and the logarithmic transformation of realized volatility when applying HAR models. Their findings produce resembling results irrespective of the employed variable specification and we will conduct replicated tests to study whether the same conclusion applies to our data set.

To further test the models' robustness we respecify the variables by applying a different starting period where we shift the forecasting periods two weeks forward in order to deviate as much as possible from the original setting. This way, we can analyze whether our results are subject to time dependency.

It is evident from Figure 1 that the non-overlapping procedure limits our potential number of forecasting periods concerning the one week forecasting horizon. The reason for the chosen specification is to facilitate a comparison between a model's forecasting accuracy across the two horizons when using the same number of forecasting periods, Δ. However, following Corsi (2009) we produce additional tests in Section 5.4 where we respecify the one week forecasting horizon, creating a series of non-overlapping forecasting periods, δ, which consist of five trading days, $RV_{t+1,t+5}$, $RV_{t+6,t+10}$, ..., $RV_{t+k,T}$. This procedure increases the number of forecasting periods to 416. Due to the respecification the one month forecasting variables, $RV_{t-21,t}$, $C_{t-21,t}$ and $J_{t-21,t}$, will experience some overlap. Section 5.4 also includes a replicated test using the original variable specification for the most actively traded exchange rate, EUR/USD, to evaluate if results are contingent on liquidity discrepancies.

3.6 Forecasting procedure

This section will cover an explanation of the implementation concerning the in-sample and out-of-sample forecasts as well as the post-test analysis. The fundamental difference between in-sample and out-of-sample tests is the structure of the data when estimating the regressions.

The in-sample procedure executes the regression using the entire sample period and the performance evaluation focuses on the adjusted R^2 as the goodness-of-fit statistic. The out-of-sample procedure on the other hand divides the sample into two periods; a fitting period $(obs_1, ..., obs_n)$, and a test period

$(obs_{n+1}, ..., obs_T)$. The model is calibrated during the fitting period while the test period is reserved to assess the model's forecasting accuracy. Selecting the length of the two periods entails a trade-off between how much data should be accounted for to calibrate the model and the length of the forecasting period to test the accuracy of the forecast.

Tashman (2000) highlights that forecasters generally agree that the out-of-sample forecasting method is the most realistic setting one can apply to evaluate a model's forecasting accuracy. Hence, the one-step ahead out-of-sample forecasts are the foundation of this paper. To be completely clear with the notation; the tests often are referred to as pseudo-out-of-sample tests since we are using historical data as the testing period, and not the actual future itself. Given the time consuming arrangement of "true" out-of-sample forecasting, where one needs to construct forecast estimates and wait until tomorrow to compute the forecasting error, the pseudo-out-of-sample setting has overwhelming time advantages.

Our employed out-of-sample technique is the rolling window estimation, a method resembling a moving average process where the oldest observation is dropped from the fitting period and the newest is added as the window moves one-step ahead, approaching time T. Thus, the fitting period is constantly updated, and according to Poon and Granger (2003) the rolling window method might be more appropriate if the model's parameter estimates exhibit non-stationarity or time variation.

For every executed one-step ahead forecast, f_{n+i}, a forecasting error is calculated, $e_{n+i} = f_{n+i} - y_{n+i}$ where y_{n+i} is the actual value. The regression model's forecasting accuracy is thereafter evaluated according to the two employed measurements, the Root Mean Squared Error (RMSE) and the Mean Average Error (MAE). As both names suggest, the RMSE is the square root of the mean squared errors and the MAE is the average of the absolute errors:

$$RMSE = \sqrt{\frac{1}{T}\sum_{n+i}^{T} e_{n+i}^2} \qquad \text{and} \qquad MAE = \frac{1}{T}\sum_{n+i}^{T} |e_{n+i}| \,. \tag{20}$$

We employ these two error measurements to evaluate the forecasting accuracy since they provide different views concerning the forecasting errors' distribution where the RMSE particularly captures and punish large outliers.

4 Data and descriptive statistics

Our empirical analysis is based upon two types of data sets collected from the Bloomberg Professional Service spanning from January 1, 2008 until February 19, 2016. We collect data on the euro (EUR) and U.S. dollar (USD) denominated in Swedish krona (SEK). According to the Triennial Central Bank Survey (2013), the USD/SEK and the EUR/SEK are the most actively traded currency pairs denominated in the Swedish krona.

The first set contains high-frequency exchange rate data regarding the aforementioned currency pairs sampled at five-minute intervals and includes information about closing and opening prices, tick count, and high and low quotes. The exchange rates are part of the Bloomberg Generic Composite (BGN), and not based on actual market trades since the contributors (e.g. regional banks, brokers, and trading platforms) do not provide trade information. Rather, BGN is an algorithm producing indications of bid and ask quotes derived from hundreds of contributors. The composite bid rate is the highest bid rate and the composite ask rate is the lowest ask rate of all of the active contributors. All contributors are evaluated for quality and consistency and for the contributors' rates to be eligible to send data to the composite they must be considered open. The algorithm therefore determines the validity of the given prices by controlling which contributors are currently open and which are closed. If the contributor is considered open, the bid is compared to all of the currently active bids from the other active providers and the highest of these bids is used to update the composite bid rate, and the composite ask rate is updated applying the same procedure. To come around the fact that contributors do not provide information about trades, the algorithm generates a trade when a best ask and a best bid is received, and the trade generated is the mid-value between these rates. Moreover, a five minute rule ensures that the quality of the data is maintained by not generating a mid-value if the best bid or best ask is more than five minutes old.

The foreign exchange market is open 24 hours a day. Bloomberg currently records currency market data between 5 p.m. ECT (10 p.m. CET) on Sunday and 5 p.m. ECT (10 p.m. CET) on Friday. Three time frames are represented on the Bloomberg in order to account for the different currency markets trading hours: Tokyo, London, and New York (Bloomberg QFX function).

In practice, the introduction of market microstructure noise due to price discreteness, bid-ask spreads, and non-synchronous trading makes it undesirable to sample data at the maximum frequency accessible. It is shown in Xu and Taylor (1997), and Andersen et al. (2001a, 2001b, 2011) that the five-minute sampling frequency is adequate and largely free of microstructure bias. The appropriate measure is therefore to make use of the five-minute sampling frequency and calculate the returns from the closing prices, which yields approximately 288 observations per day provided that trades, as defined above, actually occurred during every five-minute interval of the day. Figure 2 provides a graphical illustration of the average daily tick distribution across the sample period, suggesting that there is enough liquidity in the markets to include returns during the night. The foremost left and right parts on the two figures correspond to when the market is open in Tokyo but closed in New York and London, yielding the lowest amounts of ticks, with the average tick count exceeding 500 ticks per five-minute interval for the EUR/SEK and at least the double for the USD/SEK. The middle parts of the figures show the high peaks when the markets are open either in London or New York, or both simultaneously, but closed in Tokyo. During these peak intervals the amount of ticks is on average higher than 1000 and 2000 per five-minute interval for the EUR/SEK and USD/SEK, respectively.

Figure 2: Average daily tick count

(a) EUR/SEK (b) USD/SEK

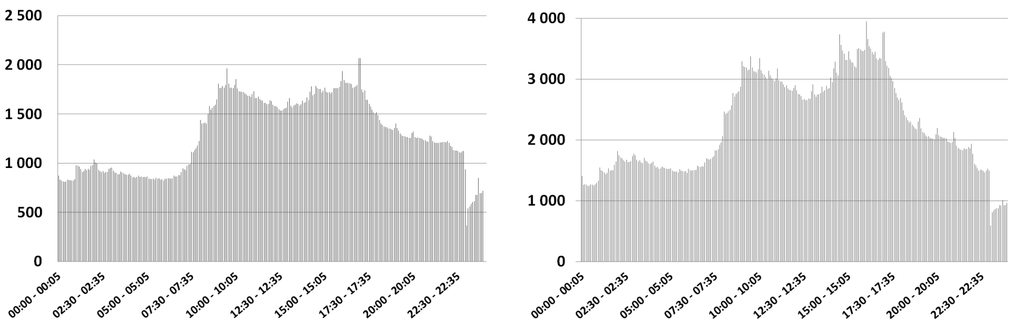

This figure provides a graphical illustration of the average daily tick distribution across the sample. The vertical axis refers to the number of ticks and the horizontal axis refers to trading hours (in CET). A single bar represents the average number of ticks in a given five minute interval.

Given that the realized volatility measure provided in equation (7) is based on the sum of the intraday squared returns it is of great importance that the number of return observations is fairly equal across the days in the sample period. If this would not be the case the undesirable effect might be that the realized volatility measures are biased due to unequal number of observations included. Therefore, we clean the data and explicitly exclude observations on Sundays (424 days during the sample period) since they usually only include 24 five-minute intervals (trading starts at 10 p.m. giving 12×2=24 five-minute intervals). Trading ends on Fridays' at 10 p.m., meaning that Fridays' are made up of less than 288 five-minute intervals (288-24=264 five-minute intervals). The minimum limit is set to 264 observations, in order to incorporate Friday observations. Hence, trading days with trades taking place in less than 264 five-minute intervals will be omitted from the data set but due to the liquidity in the foreign exchange markets this is uncommon and in Table 1 below it can be seen that the number of omitted trading days is fairly small, almost exclusively occurring around holiday periods.

Table 1: Descriptive statistics - exchange rates

	EUR/SEK	USD/SEK
Included days in sample	2 103	2 105
Omitted trading days	17	15
Number of days with jumps	2 000	1 434
No. of 5 min. intervals	599 355	599 925
Avg. No. of 5 min. intervals per day	285	285
Tick count (in millions)	775	1 355
Avg. No. ticks per 5 min interval	1 291	2 258

The table presents the characteristic for the exchange rate data sets after filtering out Sundays. *Included days in sample* refers to the trading days included from the sample. *Omitted trading days* refers to the days where trading took place in less than 264 five minute intervals. *Number of days with jumps* refers to the days where the ratio test from equation (12) indicated that a jump had taken place. *No. of 5 min intervals* is the total amount of included five-minute intervals for each currency pair. *Avg. No. of 5 min. intervals per day* is the average number of five-minute intervals, calculated as total intervals divided by included days, and the *Avg. No. ticks per 5 min. interval* is average number of ticks per five-minute intervals, calculated as the total number of ticks divided by included days.

In the model specifications, explained in Section 3.5, we use 22 trading days to note the one-month horizon, although considering public holidays and the applied filtering process does not necessarily imply 22 trading days during the lifetime of an option. The one month horizon therefore includes the trading days' average realized volatility during the option's lifetime. Applying a non-overlapping procedure across our sample period yields 97 observations to our empirical tests.

The second set of data consists of end-of-day implied volatility time series. Following the approach of Li (2002), Dunis and Huang (2002), Pong et al. (2004), Christoffersen and Mazzotta (2005), Sarantis (2006), and Kellard et al. (2010), we collect at-the-money (ATM) over-the-counter (OTC) market quoted daily implied volatilities with fixed maturities of one month. The composite is not calculated by Bloomberg but contributed as quotes in a similar fashion as described above regarding the exchange rates data[1] (Bloomberg QFX function, Bloomberg Financial Data Disclaimer). Li (2002), Christoffersen and Mazzotta (2005), and Kellard et al. (2010) highlight several advantages utilizing OTC options rather than the more common approach of backing out implied volatility from option-pricing models, using exchange traded options. First, the volumes and liquidity in the OTC market by far exceed trading on organized exchanges (BIS, Triennial Central Bank Survey 2013). To the extent that illiquidity may introduce errors in the measurements, this risk is inferior using OTC data over exchange traded options. Second, the OTC options market data allows us to match historical data with that of implied volatilities as the latter comes with constant time to expiration (from one month to six months) whereas exchange traded options' time to expiration cycle is quarterly. A third difference regard measurement errors; for exchange traded options, the issue of moneyness in the term structure may infect the relationship between the option and its underlying volatility that it is trying to capture, while the OTC options are always ATM. Lastly, the exchange traded options only trade for seven major foreign currencies, all denominated in U.S. dollars, which excludes the possibility to examine the Swedish krona.

Table 2: Descriptive statistics - regression variables

	EUR/SEK											
	$RV_{t+1,t+22}$	$RV_{t+1,t+5}$	IV_t	$RV_{t-21,t}$	$RV_{t-4,t}$	RV_t	$J_{t-21,t}$	$J_{t-4,t}$	J_t	$C_{t-21,t}$	$C_{t-4,t}$	C_t
Mean	0.090	0.090	0.081	0.089	0.088	0.089	0.045	0.041	0.042	0.079	0.078	0.077
Median	0.082	0.082	0.072	0.082	0.080	0.080	0.045	0.035	0.036	0.071	0.069	0.069
St. Dev	0.033	0.030	0.030	0.033	0.039	0.040	0.007	0.019	0.030	0.030	0.035	0.030
Min	0.044	0.046	0.043	0.044	0.046	0.044	0.024	0.016	0.000	0.039	0.040	0.037
Max	0.236	0.187	0.189	0.236	0.305	0.311	0.057	0.153	0.271	0.210	0.264	0.163
Skewness	1.916	1.219	1.879	1.922	2.626	2.482	-1.484	2.734	4.637	1.912	2.511	1.315
Kurtosis	4.502	1.502	3.621	4.485	9.728	9.789	2.417	11.544	33.398	4.178	8.659	1.322
	USD/SEK											
	$RV_{t+1,t+22}$	$RV_{t+1,t+5}$	IV_t	$RV_{t-21,t}$	$RV_{t-4,t}$	RV_t	$J_{t-21,t}$	$J_{t-4,t}$	J_t	$C_{t-21,t}$	$C_{t-4,t}$	C_t
Mean	0.140	0.141	0.133	0.140	0.137	0.132	0.045	0.044	0.036	0.132	0.128	0.123
Median	0.127	0.127	0.123	0.127	0.122	0.120	0.041	0.039	0.036	0.123	0.116	0.115
St. Dev	0.049	0.047	0.046	0.049	0.054	0.050	0.015	0.020	0.029	0.048	0.052	0.049
Min	0.066	0.068	0.067	0.066	0.064	0.057	0.025	0.018	0.000	0.061	0.059	0.050
Max	0.313	0.269	0.300	0.313	0.366	0.319	0.112	0.175	0.123	0.302	0.347	0.295
Skewness	1.571	1.194	1.422	1.572	1.971	1.341	1.948	3.403	0.662	1.509	1.797	1.280
Kurtosis	2.795	0.859	2.393	2.794	5.365	2.222	4.979	18.890	0.386	2.617	4.440	1.970

The table presents the sample characteristics (mean, median, standard deviation, minimum, maximum, skewness and kurtosis) for the variables used in the regression analysis. Mean, standard deviation, minimum, maximum, is in the form of volatility. Skewness and kurtosis measures the variables' distribution.

[1] We have been in contact with Bloomberg attempting to get further specifications on how the implied volatility measure is constructed, but they could not give us a more detailed explanation than the one provided.

The descriptive statistics for the variables included in the regressions are listed in Table 2. On average, higher values can be observed across all measures for the USD/SEK indicating that this series exhibit higher levels of volatility during our period of study. The skewness and kurtosis indicate a fat right tail distribution in the J_t variable for the EUR/SEK. The variable consists of single observations and not an average over a period which makes it exposed to outliers, considering the smoothening effect that an averaging procedure yields and that J_t lack. In Subsection 5.3 we conduct tests to see what effect such a distribution imposes on the regressed variables.

5 Results

In this section we present and analyze the results. In Subsection 5.1, the main empirical in-sample results are reported for the EUR/SEK and USD/SEK exchange rates, respectively. Subsequently, in Subsection 5.2, we provide the out-of-sample results for the same currency pairs. Thereafter, in Subsection 5.3, we put forward the analytics from the robustness tests, and finally in Subsection 5.4 we present the results from the additional tests.

5.1 In-sample estimation results

We first address the issue of analyzing the models in-sample. The first column in Table 3 and Table 4 lists the independent variables while the remaining columns report the corresponding coefficients together with the standard errors in parentheses and the adjusted R^2 ($AdjR^2$) in the bottom row.

Starting with Table 3 and EUR/SEK, the combined impact from the daily, weekly, and monthly HAR-RV components on future weekly realized volatility is 0.45+0.16+0.09=0.70, close to the first order autocorrelation of the weekly realized volatility forecast (0.67). The t-statistics are 0.77, 0.90, and 2.74, respectively, which indicates that the daily and weekly variables contain only negligible information concerning future weekly exchange rate volatility. The one month realized volatility forecast yields corresponding results, where the aggregated impact is 0.78 with a first order autocorrelation of 0.79, and the only significant variable being the $RV_{t-21,t}$ variable. Resembling results are also obtained in Table 4 for the USD/SEK, where the first order autocorrelation of the one month realized volatility forecast is 0.85 and equivalent to the combined impact from the HAR-RV variables, which is 0.85. The one week realized volatility forecasting variables combined impact (0.80) differ slightly more from the first order autocorrelation (0.73), nevertheless the results are rather close.

The combined HAR-RV-IV model yields coherent results across the two exchange rates. For the EUR/SEK one-week forecast the IV_t variable remains highly significant in the HAR-RV-IV model even though the coefficient decreases somewhat to 0.715, in contrast to when it is regressed individually (0.842), whereas none of the past realized volatility components remain significant. Indeed, the one month past realized volatility coefficient, $RV_{t-21,t}$, decreases to 0.001 from 0.449 and loses its entire explanatory power when combined with implied volatility, suggesting that IV_t is informationally efficient, subsuming all informational content from past realized volatility. An interesting finding given that IV_t by construction is a one month forward looking parameter but nevertheless seem to explain a large fraction of the short term variation. The adjusted R^2 in the bottom row of Table 3 display values in the range between 0.64 and 0.70 for the one week realized volatility forecast, and between 0.63 and 0.67 for the one month forecast while higher adjusted R^2 values are displayed in the bottom row of Table 4 for the USD/SEK, suggesting a link between high liquidity and high adjusted R^2. The adjusted R^2 values are also markedly higher than in Busch et al. (2011) considering their estimation on foreign exchange data, which might be partly linked to their use of different currencies ($/Deutsche mark) and sample period.

Table 3: EUR/SEK in-sample results

	One Week Forecast, $RV_{t+1,t+5}$					One Month Forecast, $RV_{t+1,t+22}$				
	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV
c	0.026**	0.021**	0.021**	0.003	0.011	0.020**	0.018**	0.017**	0.035**	0.040**
	(0.005)	(0.005)	(0.005)	(0.011)	(0.011)	(0.006)	(0.006)	(0.006)	(0.012)	(0.013)
RV_t	0.094	-	0.150	-	-	0.058	-	0.089	-	-
	(0.122)		(0.114)			(0.135)		(0.134)		
$RV_{t-4,t}$	0.162	-	-0.031	-	-	0.120	-	0.010	-	-
	(0.180)		(0.175)			(0.199)		(0.205)		
$RV_{t-21,t}$	0.449**	-	0.001	-	-	0.599**	-	0.345	-	-
	(0.164)		(0.191)			(0.181)		(0.224)		
IV_t	-	0.842**	0.715**	-	0.519*	-	0.878**	0.406	-	0.444
		(0.057)	(0.183)		(0.220)		(0.069)	(0.215)		(0.252)
C_t	-	-	-	0.331**	0.300**	-	-	-	0.306*	0.279*
				(0.111)	(0.109)				(0.125)	(0.125)
$C_{t-4,t}$	-	-	-	0.078	-0.101	-	-	-	-0.158	-0.311
				(0.221)	(0.229)				(0.250)	(0.262)
$C_{t-21,t}$	-	-	-	0.288	0.035	-	-	-	0.750**	0.534*
				(0.191)	(0.215)				(0.217)	(0.247)
J_t	-	-	-	-0.165	-0.088	-	-	-	-0.069	-0.003
				(0.103)	(0.106)				(0.116)	(0.121)
$J_{t-4,t}$	-	-	-	0.389	0.326	-	-	-	0.245	0.192
				(0.237)	(0.232)				(0.267)	(0.266)
$J_{t-21,t}$	-	-	-	0.511*	0.211	-	-	-	-0.490	-0.747*
				(0.258)	(0.282)				(0.292)	(0.323)
$Adj R^2$	0.640	0.688	0.688	0.688	0.703	0.630	0.628	0.640	0.665	0.673

Estimation results for the HAR-RV equation (15), IV equation (17), HAR-IV equation (18), HAR-CJ equation (16) and HAR-CJ-IV equation (19). The first column lists the independent variables while the remaining columns report the corresponding coefficients together with the standard errors in parentheses and the adjusted R^2 in the bottom row. Number of forecasting periods, Δ, is equal to 97.

** Denotes rejection at 1% significance level.

* Denotes rejection at 5% significance level.

The results for USD/SEK conform to the relationship found for the EUR/SEK currency pair, i.e. that the ex-ante implied volatility measure performs strongly over the short horizon compared to the HAR components. Forecasting over one month has the effect of altering the analysis. The one month past realized volatility measure, $RV_{t-21,t}$, is almost significant for the USD/SEK, while the opposite condition becomes true for the EUR/SEK where IV_t is instead close to being significant.

Table 4: USD/SEK in-sample results

	One Week Forecast, $RV_{t+1,t+5}$					One Month Forecast, $RV_{t+1,t+22}$				
	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV
c	0.030**	0.026**	0.027**	0.038**	0.033**	0.022**	0.020*	0.020*	0.033**	0.031**
	(0.008)	(.008)	(0.008)	(0.008)	(0.009)	(0.008)	(0.008)	(.008)	(0.009)	(.009)
RV_t	0.139	-	0.085	-	-	0.101	-	0.070	-	-
	(0.119)		(0.119)			(0.124)		(0.126)		
$RV_{t-4,t}$	0.281	-	0.239	-	-	0.210	-	0.186	-	-
	(0.171)		(0.169)			(0.179)		(0.179)		
$RV_{t-21,t}$	0.384*	-	0.129	-	-	0.544**	-	0.397	-	-
	(0.158)		(0.194)			(0.164)		(0.206)		
IV_t	-	0.866**	0.391*	-	0.348	-	0.898**	0.225	-	0.119
		(0.055)	(0.180)		(0.183)		(0.060)	(0.191)		(0.190)
C_t	-	-	-	0.226	0.172	-	-	-	0.184	0.166
				(0.129)	(0.131)				(0.133)	(0.136)
$C_{t-4,t}$	-	-	-	0.305	0.286	-	-	-	0.178	0.171
				(0.194)	(0.191)				(0.198)	(0.199)
$C_{t-21,t}$	-	-	-	0.362*	0.124	-	-	-	0.627**	0.546**
				(0.166)	(0.206)				(0.170)	(0.215)
J_t	-	-	-	-0.133	-0.145	-	-	-	-0.114	-0.119
				(0.089)	(0.088)				(0.091)	(0.092)
$J_{t-4,t}$	-	-	-	-0.053	-0.086	-	-	-	0.051	0.040
				(0.210)	(0.208)				(0.215)	(0.217)
$J_{t-21,t}$	-	-	-	-0.105	-0.093	-	-	-	-0.452	-0.449
				(0.264)	(0.260)				(0.270)	(0.271)
$Adj R^2$	0.726	0.718	0.736	0.740	0.748	0.730	0.700	0.731	0.753	0.751

Estimation results for the HAR-RV equation (15), IV equation (17), HAR-IV equation (18), HAR-CJ equation (16) and HAR-CJ-IV equation (19). The first column lists the independent variables while the remaining columns report the corresponding coefficients together with the standard errors in parentheses and the adjusted R^2 in the bottom row. Number of forecasting periods, Δ, is equal to 97.
** Denotes rejection at 1% significance level.
* Denotes rejection at 5% significance level.

Hence, the $RV_{t-21,t}$ and IV_t variables both become insignificant when combined in the HAR-RV-IV model despite being highly significant in the HAR-RV and IV model, respectively. The fact that the adjusted R^2 remains high raises concerns about multicollinearity. Figure 3 provide graphical illustrations over the high correlation among the ex-ante variables, IV_t, and the ex-post variables, $RV_{t-21,t}$, for both currencies where it is evident that the two measures follow each other surprisingly well, albeit some small deviations can be distinguished over the sample period.

Figure 3: Relation between IV_t and $RV_{t-21,t}$

(a) EUR/SEK (b) USD/SEK

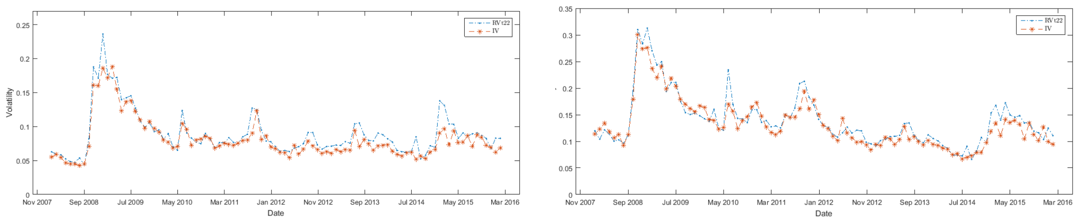

The figures provide graphical illustrations of the high correlation between IV_t and $RV_{t-21,t}$. Plot (a) illustrates the correlation for EUR/SEK and Plot (b) illustrates the correlation for USD/SEK.

Supplementary numerical tests for correlation warrant the concern and it becomes evident from Table 5 that IV_t and $RV_{t-21,t}$ exhibits high correlation where the correlation coefficients for the $RV_{t-21,t}$ and IV_t variables are approximately 0.95 for both currency pairs.

Table 5: Correlation matrices

	EUR/SEK					USD/SEK			
	RV_t	$RV_{t-4,t}$	$RV_{t-21,t}$	IV_t		RV_t	$RV_{t-4,t}$	$RV_{t-21,t}$	IV_t
RV_t	1.000				RV_t	1.000			
$RV_{t-4,t}$	0.923	1.000			$RV_{t-4,t}$	0.907	1.000		
$RV_{t-21,t}$	0.865	0.940	1.000		$RV_{t-21,t}$	0.863	0.947	1.000	
IV_t	0.831	0.918	0.944	1.000	IV_t	0.869	0.925	0.949	1.000

The left correlation matrix correspond to the independent variables included in the HAR-RV-IV model (6) for the EUR/SEK. The right correlation matrix correspond to the independent variables included in the HAR-RV-IV model (6) for the USD/SEK.

In order to make further tests for multicollinearity we drop the $RV_{t-21,t}$ from the one month in-sample HAR-RV-IV model and regress the $RV_{t+1,t+22}$ one month forecast on the RV_t, $RV_{t-4,t}$, and IV_t variables. One can observe, looking at Table 6 below, that the outcome highly resembles the original outcome from the HAR-RV model for the USD/SEK. On the other hand, for the EUR/SEK, the $RV_{t-4,t}$ becomes significant in addition to the IV_t variable that works as a substitute for the $RV_{t-21,t}$ variable from the original model, generating slightly different results in comparison with the original HAR-RV model.

Table 6: One month forecast, $RV_{t+1,t+22}$

	EUR/SEK	USD/SEK
c	0.020**	0.024**
	(0.006)	(0.008)
RV_t	0.101	0.045
	(0.135)	(0.127)
$RV_{t-4,t}$	0.134	0.372**
	(0.190)	(0.153)
IV_t	0.605**	0.447**
	(0.173)	(0.151)
$Adj R^2$	0.646	0.732

In-sample estimation results where the $RV_{t-21,t}$ variable is dropped and replaced by IV in order to evaluate multicollinearity. $\Delta = 97$
** Denotes rejection at 1% significance level.

Apart from raising econometric concerns, the fact that IV_t and $RV_{t-21,t}$ are highly correlated raises interesting questions about the IV_t measure itself. As explained in Section 4, the information provided by Bloomberg regarding the IV_t measure is not entirely unambiguous. Regardless of the reasons for the high correlation, be it due to randomness, Bloombergs Professional Services algorithm, or if the pricing of options heavily rely on past volatility, one can question how much of an ex-ante measure this implied volatility is.

Progressing on to analyze the HAR-CJ and HAR-CJ-IV models in Table 3 and Table 4 we find that the separation of the continuous sample path, C, and jump, J, does not change the adjusted R^2 substantially in comparison with the original HAR-RV model. The adjusted R^2 is only slightly higher for HAR-CJ than HAR-RV in all regressions, whereby this finding conforms to previous studies (see Andersen et al. [2007] and Busch et al. [2011]). The one week-forecast for the EUR/SEK in Table 3 shows that IV_t appears to be informationally efficient over the heterogeneous autoregressive components in the clustered HAR-CJ-IV model, where only one variable, C_t, remains significant. Extending the forecasting horizon to one month yields a decreased significance level for the IV_t, where also half of the autoregressive components turn significant. Forecasting on the one week horizon for the USD/SEK in Table 4 provides a significant $C_{t-21,t}$ component in the HAR-CJ model, although it drops completely after the inclusion of IV_t in the HAR-CJ-IV model where none of the variables are significant. These results signal that the multicollinearity issue appears to be present in this scenario as well. Shifting focus yet again to the one month forecasting horizon shows that $C_{t-21,t}$ is the only significant variable in both the HAR-CJ and HAR-CJ-IV models, with higher adjusted R^2 in the HAR-CJ setup, which designates that IV_t has a weak predictability in this setting.

5.2 Out-of-sample estimation results

Table 7 and Table 8 summarize the models' out-of-sample forecasting performance for the EUR/SEK and the USD/SEK, respectively. As described in Section 3.6, we adopt two different evaluation measures that consider the forecasting errors; the root mean square error and the mean absolute error. The tables report the RMSE and MAE one step-ahead forecasting errors for the rolling window estimates with a fitting period length of 66 observations. The overall analysis is that there are no huge deviations in terms of error statistics across the models; although some clear patterns are present in the test statistics.

Table 7: EUR/SEK out-of-sample results

	One Week Forecast, $RV_{t+1,t+5}$					One Month Forecast, $RV_{t+1,t+22}$				
	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV
RMSE	**1.752**	1.818	1.804	1.851	1.816	1.574	1.632	**1.570**	1.590	1.571
MAE	**1.170**	1.214	1.273	1.261	1.249	1.044	1.107	**1.012**	1.072	1.063

Estimation results for the HAR-RV equation (15), IV equation (17), HAR-IV equation (18), HAR-CJ equation (16) and HAR-CJ-IV equation (19). The table reports the RMSE and MAE one step-ahead forecasting errors (× 100) for the rolling window estimates, with a fitting period length of 66. The bold text shows the model with the lowest RMSE and MAE over the specific horizon.

Taking Table 7 under consideration for the EUR/SEK exchange rate it shows that the HAR-RV model performs best over the one week horizon. The second best performance, resting on to the RMSE measure, is achieved by the HAR-RV-IV model. Consequently, since the HAR-RV-IV model performs better than the IV model, it indicates that the relatively good forecasting performance of the HAR-RV-IV model should be ascribed to the HAR components rather than IV. The same outcome can be observed studying the one month forecast where HAR-RV-IV exhibits the lowest forecasting errors and the IV model the highest according to both measures, suggesting that a substantial fraction of the favorable performance in the HAR-RV-IV can be linked to the HAR-RV model. The model ranking is consistent between the two error measurements for the one month forecasting horizon, meanwhile some deviations can be noticed for the one week forecasts. In that instance, the IV model is ranked second best according to the MAE measure, while it is ranked at the bottom according to the RMSE measure, and considering that RMSE penalize outliers indicates that the IV model might exhibit some relatively large outliers forecasting the one week horizon.

Table 8: USD/SEK out-of-sample results

	One Week Forecast, $RV_{t+1,t+5}$					One Month Forecast, $RV_{t+1,t+22}$				
	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV
RMSE	2.027	**1.856**	1.900	2.032	2.036	1.789	1.865	**1.788**	1.908	1.969
MAE	1.437	**1.297**	1.314	1.387	1.338	1.468	1.378	1.406	1.363	**1.345**

Estimation results for the HAR-RV equation (15), IV equation (17), HAR-IV equation (18), HAR-CJ equation (16) and HAR-CJ-IV equation (19). The table reports the RMSE and MAE one step-ahead forecasting errors (× 100) for the rolling window estimates, with a fitting period length of 66. The bold text shows the model with the lowest RMSE and MAE over the specific horizon.

Next, looking at Table 8 for the USD/SEK it is evident that there are some inconsistencies in the out-of-sample tests amid the two exchange rates, as was the case for the in-sample results. For instance, all models exhibit higher forecasting errors, and the IV model is superior to the other models when forecasting one week volatility, in regards to both the RMSE and MAE measures, with HAR-RV-IV being the runner-up model. Moreover, the one month forecasts display big differences in the model ranking, in contrast to the one month forecasts for the EUR/SEK in Table 7. The HAR-CJ-IV and HAR-CJ models drop from best and runner-up to last and second to last with reference to RMSE instead of MAE, indicating that the models possess large outliers.

Comparing the results from the out-of-sample tests with the in-sample findings entail interesting remarks. For instance, a low correlation is recognized between a models ability to explain in-sample variation with its out-of-sample forecasting accuracy. The reasoning can primarily be applied to the more sophisticated HAR-CJ models that performs relatively poor out-of-sample forecasts but yields high in-sample adjusted R^2, suggesting that the extended variable modeling is not worth-while. The findings are in line with the results presented in Busch et al. (2011).

5.3 Robustness tests

Additional tests are conducted in order to verify the robustness of the previously reported results, with delicate changes to the models' forecasting performance and their characteristics. This is done by examining whether our results are dependent on variable transformation, and also if time is an influencing factor.

We begin with respecifying the variables into realized variance and logarithmic form. We find that many previous studies that conduct volatility forecasting specify the models in realized variance form, although Andersen et al. (2007) find resembling results using realized variance, realized volatility or the logarithm of realized volatility when applying HAR models. The left side of Panel A in Table 9, included in the Appendix, refers to the in-sample tests for the one week horizon for the EUR/SEK with the variables taking the variance form. The results fall in line with the original specifications found in Table 3 for the HAR-RV, IV and HAR-IV models with only minor differences with reference to the coefficients. The C variables in the HAR-CJ and HAR-CJ-IV models are similar to the results in Table 3 and 4, however, the results for the jump components are deviating. Further analysis indicate that the large positive coefficients for $J_{t-4,t}$ and $J_{t-21,t}$ in Table 9 induce that the presence of jumps over the past week and month imply higher future volatility. The large statistical significance is surprising and we analyze this deviation from the original results as a weakness of the model. The out-of-sample test, displayed on the left hand side of Table 9, Panel B, indicates consistency since the HAR-RV model performs best with the HAR-RV-IV as the runner up, equivalent to the findings in the original setting.

Figure 4: Distribution graphs

(a) Original form (b) Logarithmic transformation

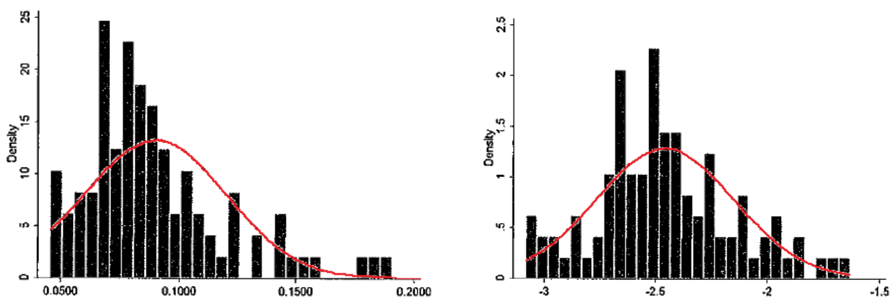

Figure (a) shows the distribution of the weekly $RV_{t+1,t+5}$ variable for the EUR/SEK defined in its original form. Figure (b) shows the distribution for the same variable after the logarithmic transformation.

Table 2 highlighted the presence of large skewness and kurtosis in some of the variables. To test the effect this imposes on the results we conduct a one week in-sample forecast on EUR/SEK with the variables respecified into logarithmic form. The change in distribution is observable in Figure 4 following the transformation to logarithmic form for the $RV_{t+1,t+5}$ variable. The logarithmic form in (b) takes care of some of the fat right tail in (a), as predicted. Subsequently, the right hand side of Panel A in Table 9 reports the results from the in-sample estimations. Clearly the coefficients change in accordance to the respecification but the overall picture resembles the original test in Table 3, indicating that the results are not sensitive to the manner in which realized volatility is specified. The out-of-sample test results' in Panel B indicate that the respecification favor the IV model where it performs best followed by the HAR-CJ-IV model. The improved performance of the HAR-CJ and HAR-CJ-IV models is anticipated given the large skewness and kurtosis in the J variables in the original setting.

To further check for robustness we perform tests using the original volatility specification, but apply a different starting period when creating our variables. To deviate as much as possible from the original setting we move the starting point two weeks forward. If our results are robust they should not deviate due to this change since it would indicate that there is a time dependency. A consequence of this respecification is that we lose one forecasting period, Δ, leaving us with 96 periods. Figure (a) in Figure 7 and Figure 8 in the Appendix illustrates that the one week robust horizon frequently deviates from the original specification, whereas Figure (b) in Figure 7 and Figure 8 shows that the one month robust horizon follows the original horizon to a much greater extent. The reason might be that the one month horizon incorporates overlapping observations since half of the original one month horizon overlaps with half of the robust one month forecasting horizon. Meanwhile, the one week horizon is completely non-overlapping.

Panel A in Table 10 and 11 in the Appendix display the results for the EUR/SEK and USD/SEK robustness estimations. There are clear similarities concerning the in-sample results between the two exchange rates in this setting. However, the results deviate significantly from the original tests. For both exchange rates and across both forecasting horizons the RV_t variables are significant in the HAR-RV and HAR-RV-IV models, a deviation from previous results. Hence, we can no longer confirm that the implied volatility measure is informationally efficient in the one week horizon, rather the HAR components are the only significant variables in the combined models. The adjusted R^2 display large similarities with the original tests, with the exception that it is higher for the HAR-RV model than the HAR-CJ for the EUR/SEK.

The EUR/SEK out-of-sample results in Table 10, Panel B, for the one week forecasting horizon is significantly higher than in any other test. This can be explained by the large forecasting errors around the year-end of 2014, illustrated for the HAR-RV model in Figure 5, as well as the discrepancy between the one week horizons described above. The error statistics are similar across all models but with the HAR-CJ-IV displaying the lowest forecasting error in contrast to its relatively poor out-of-sample performances in the original setting. This would suggest that modeling for jumps through the bipower and tripower separation of realized volatility is beneficial in volatile settings. However, apart from the lastly ascribed scenario, the overall pattern indicated in the original tests seems to hold where the simpler HAR, IV, and HAR-IV models perform better out-of-sample forecasts.

Figure 5: EUR/SEK: The HAR-RV model's forecast of $RV_{t+1,t+5}$

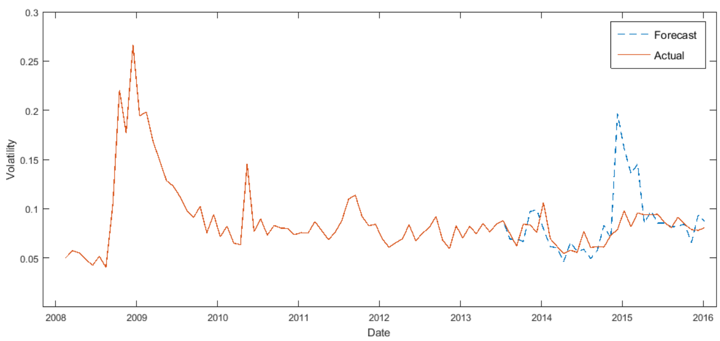

The HAR-RV model's performance in forecasting the EUR/SEK one week forecast, $RV_{t+1,t+5}$.

In summary, variable transformation only has minor influence on the in-sample and out-of-sample results. Meanwhile, applying a different period of origin when creating the variables influence the in-sample results, and this probably relates to the fairly low number of forecasting periods in the original setting. We perform additional tests in the subsequent section to address the issue.

5.4 Additional tests

This section includes analysis of two different types of test. The first concerns the respecification of the one week forecasts explained in Subsection 3.5 to increase the number of forecasting periods whereas the other is a replicated test for the EUR/USD exchange rate using the original variable setting.

In Table 12 in the Appendix the results for the respecification of the one week forecasts are presented. The left hand side of Panel A includes the in-sample results for the EUR/SEK, and the right hand side displays the results for the USD/SEK. Significantly higher adjusted R^2 are evident for the HAR components, but only slightly higher for the IV measure, compared to the original results. It is related to the intensified autocorrelation in the $RV_{t+1,t+5}$ variables, displayed in Figure 6, which helps to amplify the goodness-of-fit statistic for the in-sample regressions for the HAR models. All components of the HAR-RV models are now statistically significant at one percent significance level, in contrast to the previous case found in Table 3 and 4 where only the $RV_{t-21,t}$ variable was significant.

Figure 6: Sample autocorrelation function

(a) EUR/SEK

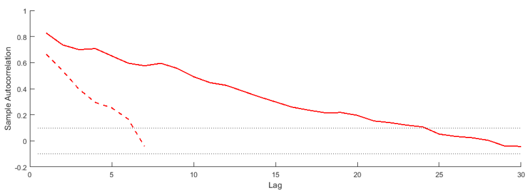

(b) USD/SEK

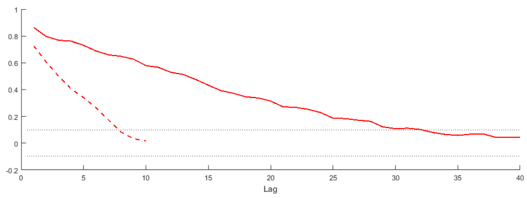

Comparisons between the $RV_{t+1,t+5}$ variable (dashed) in the original period specification Δ and the $RV_{t+1,t+5}$ variable (solid) in the respecified period specification δ.

A consequence of the respecification is that the IV variables are sampled the day before every new forecasting period, consisting of five trading days, and not once every month as was the case in the original setting. Given that the IV measure is based on options with one-month time to expiration, the IV observations will experience overlapping in its forecasting horizon. As mentioned previously, earlier studies suggest that this specification should be avoided, nevertheless the adjusted R^2 becomes higher for the IV models in this setting in both series compared to the original results, even if the increase is marginal. Moreover, the previously found relationship in the original one week forecasts where the $RV_{t-21,t}$ and $C_{t-21,t}$ loses significance when IV is included is still present.

The out-of-sample results in Table 12, Panel B, show some deviations from the original tests. The most conspicuous is the improved performance of the HAR-CJ-IV model across both series and the rather weak forecasting accuracy assigned to the HAR-RV model. Hence, our previous argument advocating for the use of the simpler models when performing out-of-sample forecasts does not apply to this setting with an extended number of forecasting periods.

Finally, Table 13 presents the results from the replicated test on the EUR/USD exchange rate, displaying similar results across both the in-sample and out-of-sample forecasts. The higher adjusted R^2 fall in line to the previous findings in Subsection 5.1 where the in-sample results for the USD/SEK provided higher adjusted R^2 than what was found for the EUR/SEK, indicating that higher liquidity might be linked to higher adjusted R^2. Over both forecasting horizons the modeling of C and J provides only marginal gains in the adjusted R^2 values and the implied volatility does not appear to exhibit any incremental information over the autoregressive components in any of the combined models. The out-of-sample results confirm this relationship where the simple HAR-RV model outperforms all other models across both horizons, independent of the employed error statistic. In contrast, the IV model has the highest forecasting error regardless of error statistic and forecasting horizon. Thus, in contrast to the preceding test, the results from the EUR/USD reinforce our opinion that simple models outperform more complex models when conducting out-of-sample forecasts.

6 Conclusion

This paper examines a set of models' forecasting accuracy of realized volatility in two SEK denominated exchange rates, EUR/SEK and USD/SEK, with the purpose to evaluate whether one should rely on ex-post or ex-ante measures when conducting volatility forecasts. The ex-post forecasts are executed using two types of Heterogeneous Autoregressive (HAR) models based on high-frequency data. We employ the simple Heterogeneous Autoregressive model of Realized Volatility, HAR-RV, proposed by Corsi (2009) and a modified version using the bipower and tripower variation to separate the continuous sample path (C) and the jump (J) component of realized volatility, HAR-CJ, proposed in Andersen et al. (2007). The ex-ante forecasting model is constructed by daily data on OTC implied volatility (IV). Moreover, two combined models, HAR-RV-IV and HAR-CJ-IV, are employed in order to evaluate whether a combination of ex-post and ex-ante measures strengthen the forecasting accuracy as well as to see if any of the measures are informationally efficient over the other. The forecasting performance is analyzed based on results from in-sample and out-of-sample forecasts.

The results indicate a low correlation between in-sample fit and out-of-sample forecasting performance where the HAR-CJ and HAR-CJ-IV models produce marginally higher in-sample adjusted R^2 but provide relatively poor out-of-sample forecasts, suggesting that simpler models are favorable. This argument is on the other hand contradicted when we respecify the one-week forecasting horizon with the purpose to increase the number of forecasting periods, since the HAR-CJ-IV outperforms the other models in this setting. Caution should therefore be taken when advocating simple models as the results might depend on the employed methodological approach. Our findings cannot confirm whether ex-post or ex-ante forecasting models provide the most accurate forecast since the results deviate between the currency pairs and the employed methodological approach. However, across all tests the combined HAR-RV-IV model performs well, indicating that this combined ex-post and ex-ante model might be the golden mean, and perhaps the watershed should not concern whether to rely on ex-post or ex-ante forecasting models but rather how these can be combined in the most efficient way.

7 References

Andersen, T. G. & Bollerslev, T. (1998). *Deutschemark-dollar volatility: Intraday activity patterns, macroeconomic announcements, and longer run dependencies*, Journal of Finance, 53(1), 219–266.

Andersen, T. G., Bollerslev, T., Diebold, F. X. & Ebens, H. (2001a). *The distribution of realized stock return volatility*, Journal of Financial Economics, 61, 43–76.

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

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

Andersen, T. G., Bollerslev, T. & Meddahi, N. (2011). *Realized volatility forecasting and market microstructure noise*, Journal of Econometrics, 160(1), 220-234.

Barndorff-Nielsen, O.E. & Shephard, N. (2002a). *Estimating quadratic variation using realized variance*, Journal of Applied Econometrics, 17, 457–477.

Barndorff-Nielsen, O.E. & Shephard, N. (2002b). *Econometric analysis of realized volatility and its use in estimating stochastic volatility models*, Journal of the Royal Statistical Society: Series B (Statistical Methodology), 64(2), 253-280.

Barndorff-Nielsen, O.E. & Shephard, N. (2004). *Power and Bipower Variation with Stochastic Volatility and Jumps*, Journal of Financial Econometrics, 2(1), 1-37.

Barndorff-Nielsen, O.E., Shephard, N. & Winkel, M. (2006). *Limit theorems for multipower variation in the presence of jumps*, Stochastic Processes and their Applications, 116, 796–806.

Barndorff-Nielsen, O.E. & Shephard, N. (2007). *Variation, jumps, market frictions, and high frequency data in financial econometrics. In: Blundell, R., Persson, T., Newey, W.K. (Eds.), Advances in Economics and Econometrics:*

Theory and Applications, Ninth World Congress. Cambridge University Press, 328–372.

Black, F. & Scholes, M. (1973). *The pricing of options and corporate liabilities*, The Journal of Political Economy, 637-654.

Blair, B. J., Poon, S-H. & Taylor, S.J. (2001). *Forecasting S&P 100 volatility: the incremental information content of implied volatilities and high-frequency index returns*, Journal of Econometrics, 105(1), 5-26.

Bloomberg Financial Data Disclaimers: http://www.bloomberg.com/financial-data (2016.05.01, 12:14:13 CET).

Bollerslev, T. (1986). *Generalized Autoregressive Conditional Heteroskedasticity*, Journal of Econometrics, 31, 307–28.

Busch, T., Christensen, B. J. & Nielsen, M. Ø. (2011). *The role of implied volatility in forecasting future realized volatility and jumps in foreign exchange, stock, and bond markets*, Journal of Econometrics, 160(1), 48-57.

Canina, L. & Figlewski, S. (1993). *The informational content of implied volatility*, Review of Financial Studies, 6(3), 659-681.

Christensen, B. J. & Prabhala, N. R. (1998). *The relation between implied and realized volatility*, Journal of Financial Economics, 50(2), 125-150.

Christoffersen, P. & Mazzotta, S. (2005). *The Accuracy of Density Forecasts from Foreign Exchange Options*, Journal of Financial Econometrics, 3(4), 578-605.

Corsi, F. (2009). *A simple approximate long memory model of realized volatility*, Journal of Financial Econometrics, 7(2), 174–196.

Day, T. E. & Lewis, C. M. (1992). *Stock market volatility and the information content of stock index options*, Journal of Econometrics, 52(1), 267-287.

Dunis, C. L. & Huang, X. (2002). *Forecasting and trading currency volatility: An application of recurrent neural regression and model combination*, Journal of forecasting, 21(5), 317-354.

Engle, R. F. (1982). *Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation*, Econometrica, 50(4), 987–1007.

Fleming, J. (1998). *The quality of market volatility forecasts implied by S&P 100 index option prices*, Journal of Empirical Finance, 5(4), 317–345.

Hansen, P. R. & Lunde, A. (2006). *Realized variance and market microstructure noise (with discussion).* Journal of Business and Economic Statistics, 24, 127–161.

Hull, J. C. (2014). *Options, futures, and other derivatives.* Upper Saddle River, N.J., Pearson/Prentice Hall.

Jorion, P. (1995). *Predicting volatility in the foreign exchange market*, The Journal of Finance, 50(2), 507-528.

Kellard, N., Dunis, C. & Sarantis, N. (2010). *Foreign exchange, fractional cointegration and the implied–realized volatility relation*, Journal of Banking & Finance, 34(4), 882-891.

Lamoureux, C. G. & Lastrapes, W. D. (1993). *Forecasting stock-return variance: Toward an understanding of stochastic implied volatilities*, Review of Financial Studies, 6(2), 293-326.

Li, K. (2002). *Long-Memory versus Option-Implied Volatility Prediction*, Journal of Derivatives, 9(3), 9–25.

Liu, L. Y., Patton, A. J. & Sheppard, K. (2015). *Does anything beat 5-minute RV? A comparison of realized measures across multiple asset classes.* Journal of Econometrics, 187(1), 293-311.

Martens, M. & Zein, J. (2004). *Predicting financial volatility: high-frequency time-series forecasts vis-à-vis implied volatility*, Journal of Futures Markets, 24(11), 1005–1028.

Merton, R. C. (1973). *Theory of Rational Option Pricing*, The Bell Journal of Economics and Management Science, 4(1), 141-183.

Müller, U. A., Dacorogna, M. M., Nagler, R. J., Olsen, R. B. & Pictet, O. V. (1993). *A geographical model for the daily and weekly seasonal volatility in the foreign exchange market*, Journal of International Money and Finance, 12(4), 413-438.

Pong, S., Shackleton, M. B., Taylor, S. J. & Xu, X. (2004). *Forecasting currency volatility: a comparison of implied volatilities and AR(FI)MA models*, Journal of Banking & Finance, 28(10), 2541–2563.

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

Sarantis, N. (2006). *Testing the uncovered interest parity using traded volatility, a time-varying risk premium and heterogeneous expectations*, Journal of International Money and Finance, 25(7), 1168-1186.

Scott, E. & Tucker, A. L. (1989). *Predicting currency return volatility*, Journal of Banking & Finance, 13(6), 839-851.

Tashman, L. J. (2000). *Out-of-sample tests of forecasting accuracy: an analysis and review*, International Journal of Forecasting, 16(4), 437-450.

Triennial Central Bank Survey of foreign exchange and derivatives market activity in 2013: http://www.bis.org/publ/rpfx13.htm (2016.04.23, 10:45:19 CET)

Xu, X. & Taylor, S. J. (1995). *Conditional volatility and the informational efficiency of the PHLX currency options market*, Journal of Banking & Finance, 19(5), 803-821.

Xu, X. & Taylor, S. J. (1997). *The incremental volatility information in one million foreign exchange quotations*, Journal of Empirical Finance, 4, 317-340.

8 Appendix

Table 9: EUR/SEK variable respecification results

Panel A: In-sample estimation results

	Variance form, One Week Forecast, $RV_{t+1,t+5}$					**Logarithmic form**, One Week Forecast, $RV_{t+1,t+5}$				
	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV
c	0.00**	0.000**	0.000**	0.003	0.000	0.024**	0.273**	0.206**	0.271**	0.250**
	(0.000)	(0.000)	(0.000)	(0.011)	(0.000)	(0.005)	(0.006)	(0.044)	(0.028)	(0.028)
RV_t	0.014	-	0.158	-	-	0.101	-	0.145	-	-
	(0.133)		(0.117)			(0.120)		(0.112)		
$RV_{t-4,t}$	-0.080	-	-0.177	-	-	0.195	-	0.057	-	-
	(0.178)		(0.155)			(0.180)		(0.170)		
$RV_{t-21,t}$	0.714**	-	-0.074	-	-	0.420**	-	0.004	-	-
	(0.156)		(0.194)			(0.164)		(0.183)		
IV_t	-	0.839**	0.974**	-	0.644**	-	0.878**	0.054**	-	0.053**
		(0.057)	(0.172)		(0.205)		(0.069)	(0.013)		(0.019)
C_t	-	-	-	0.363**	0.326**	-	-	-	0.015	0.014
				(0.129)	(0.124)				(0.009)	(0.008)
$C_{t-4,t}$	-	-	-	0.097	-0.140	-	-	-	0.017	-0.001
				(0.229)	(0.231)				(0.017)	(0.018)
$C_{t-21,t}$	-	-	-	0.222	-0.122	-	-	-	0.025	0.002
				(0.211)	(0.229)				(0.016)	(0.017)
J_t	-	-	-	-0.447**	-0.253	-	-	-	-0.000	-0.001
				(0.139)	(0.146)				(0.006)	(0.006)
$J_{t-4,t}$	-	-	-	1.302**	1.150*	-	-	-	0.001	0.007
				(0.478)	(0.459)				(0.010)	(0.008)
$J_{t-21,t}$	-	-	-	1.478*	0.557	-	-	-	0.001	-0.012*
				(0.660)	(0.695)				(0.01)	(0.011)
$Adj R^2$	0.593	0.692	0.699	0.750	0.750	0.646	0.688	0.697	0.646	0.673

Panel B: Out-of-sample estimation results

	Variance form, One Week Forecast, $RV_{t+1,t+5}$					**Logarithmic form**, One Week Forecast, $RV_{t+1,t+5}$				
	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV
RMSE	**1.260**	1.390	1.320	1.380	1.340	1.992	**1.929**	1.976	1.989	1.953
MAE	**0.797**	0.853	0.866	0.819	0.837	1.407	**1.368**	1.431	1.401	1.373

Estimation results for the realized variance models $X^2_{t+1,t+h} = \beta_0 + \beta_d X^2_t + \beta_w X^2_{t-4,t} + \beta_m X^2_{t-21,t} + \epsilon_{t+1,t+h}$ and the logarithmic realized variance models $\log(X_{t+1,t+h}) = \beta_0 + \beta_d \log(X_t) + \beta_w \log(X_{t-4,t}) + \beta_m \log(X_{t-21,t}) + \epsilon_{t+1,t+h}$, where $X = \{RV, IV, C, J\}$. Panel A presents the in-sample regression estimates; the first column lists the independent variables while the remaining columns report the corresponding coefficients together with their standard errors in parentheses and the adjusted R^2 in the bottom row. Panel B reports the average RMSE and MAE one step-ahead forecasting errors ($\times$100000 in variance form and $\times$100 in logarithmic form) for the rolling window estimates, with a fitting period size of 66. Text in bold shows the model with lowest RMSE and MAE over the specific horizon. Number of forecasting periods, Δ, is equal to 97.

** Denotes rejection at 1% significance level.

* Denotes rejection at 5% significance level.

Table 10: EUR/SEK robustness results

Panel A: In-sample estimation results

	One Week Forecast, $RV_{t+1,t+5}$					One Month Forecast, $RV_{t+1,t+22}$				
	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV
c	0.002 (0.007)	0.002 (0.008)	0.001 (0.007)	0.003 (0.008)	0.002 (0.008)	0.010 (0.006)	0.012* (0.006)	0.009 (0.006)	0.010 (0.006)	0.009 (0.007)
RV_t	0.743** (0.155)	-	0.726** (0.160)	-	-	0.439** (0.125)	-	0.393** (0.128)	-	-
$RV_{t-4,t}$	-0.164 (0.225)	-	-0.184 (0.231)	-	-	0.186 (0.182)	-	0.132 (0.184)	-	-
$RV_{t-21,t}$	0.454** (0.148)	-	0.389 (0.211)	-	-	0.305** (0.119)	-	0.135 (0.168)	-	-
IV_t	-	1.088** (0.089)	0.118 (0.270)	-	0.082 (0.287)	-	0.960** (0.068)	0.308 (0.215)	-	0.281 (0.229)
C_t	-	-	-	0.751** (0.176)	0.743** (0.180)	-	-	-	0.401** (0.142)	0.373* (0.143)
$C_{t-4,t}$	-	-	-	-0.259 (0.176)	-0.269 (0.273)	-	-	-	0.199 (0.217)	0.167 (0.218)
$C_{t-21,t}$	-	-	-	0.484 (0.248)	0.432 (0.310)	-	-	-	0.292 (0.199)	0.117 (0.247)
J_t	-	-	-	0.201 (0.175)	0.189 (0.181)	-	-	-	0.199 (0.141)	0.159 (0.144)
$J_{t-4,t}$	-	-	-	0.062 (0.346)	0.058 (0.348)	-	-	-	-0.058 (0.278)	-0.070 (0.277)
$J_{t-21,t}$	-	-	-	0.084 (0.410)	0.084 (0.412)	-	-	-	0.147 (0.329)	0.146 (0.328)
$AdjR^2$	0.696	0.610	0.693	0.687	0.684	0.720	0.675	0.693	0.713	0.715

Panel B: Out-of-sample estimation results

	One Week Forecast, $RV_{t+1,t+5}$					One Month Forecast, $RV_{t+1,t+22}$				
	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV
RMSE	3.000	3.046	3.020	3.020	**2.982**	**1.620**	1.634	1.637	1.658	1.689
MAE	1.773	1.836	1.763	1.729	**1.694**	**1.080**	1.027	1.064	1.077	1.102

Estimation results for the HAR-RV equation (15), IV equation (17), HAR-IV equation (18), HAR-CJ equation (16) and HAR-CJ-IV equation (19). Panel A presents the in-sample regression estimates; the first column lists the independent variables while the remaining columns report the corresponding coefficients together with their standard errors in parentheses and the adjusted R^2 in the bottom row. Panel B reports the average RMSE and MAE one step-ahead forecasting errors ($\times 100$) for the rolling window estimates, with a fitting period length of 66. Text in bold shows the model with lowest RMSE and MAE over the specific horizon. Number of forecasting periods, Δ, is equal to 96.

** Denotes rejection at 1% significance level.

* Denotes rejection at 5% significance level.

Table 11: USD/SEK robustness results

Panel A: In-sample estimation results

	One Week Forecast, $RV_{t+1,t+5}$					One Month Forecast, $RV_{t+1,t+22}$				
	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV
c	-0.007 (0.007)	-0.006 (0.009)	-0.010 (0.008)	-0.004 (0.008)	-0.006 (0.009)	0.012 (0.008)	0.013 (0.009)	0.009 (0.008)	0.016 (0.009)	0.014 (0.009)
RV_t	0.401** (0.096)	-	0.366** (0.100)	-	-	0.303** (0.098)	-	0.265** (0.102)	-	-
$RV_{t-4,t}$	0.117 (0.152)	-	0.120 (0.151)	-	-	0.147 (0.156)	-	0.151 (0.155)	-	-
$RV_{t-21,t}$	0.532** (0.111)	-	0.396** (0.152)	-	-	0.471** (0.114)	-	0.328* (0.153)	-	-
IV_t	-	1.101** (0.065)	0.200 (0.155)	-	0.175 (0.160)	-	0.963** (0.063)	0.211 (0.159)	-	0.163 (0.164)
C_t	-	-	-	0.401** (0.113)	0.377** (0.115)	-	-	-	0.318** (0.116)	0.296* (0.118)
$C_{t-4,t}$	-	-	-	0.089 (0.208)	0.071 (0.208)	-	-	-	0.140 (0.212)	0.123 (0.212)
$C_{t-21,t}$	-	-	-	0.557** (0.173)	0.454* (0.197)	-	-	-	0.483 (0.176)	0.387 (0.201)
J_t	-	-	-	0.129 (0.091)	0.118 (0.091)	-	-	-	0.011 (0.092)	0.001 (0.093)
$J_{t-4,t}$	-	-	-	0.063 (0.190)	0.095 (0.192)	-	-	-	0.016 (0.193)	0.046 (0.196)
$J_{t-21,t}$	-	-	-	-0.049* (0.309)	-0.095 (0.312)	-	-	-	-0.024 (0.315)	-0.067 (0.318)
$AdjR^2$	0.835	0.751	0.838	0.841	0.843	0.786	0.716	0.790	0.795	0.798

Panel B: Out-of-sample estimation results

	One Week Forecast, $RV_{t+1,t+5}$					One Month Forecast, $RV_{t+1,t+22}$				
	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV
RMSE	1.622	**1.421**	1.572	1.635	1.636	2.137	**1.984**	2.151	2.160	2.202
MAE	1.251	**1.155**	1.217	1.239	1.226	1.579	**1.500**	1.649	1.544	1.594

Estimation results for the HAR-RV equation (15), IV equation (17), HAR-IV equation (18), HAR-CJ equation (16) and HAR-CJ-IV equation (19). Panel A presents the in-sample regression estimates; the first column lists the independent variables while the remaining columns report the corresponding coefficients together with their standard errors in parentheses and the adjusted R^2 in the bottom row. Panel B reports the average RMSE and MAE one step-ahead forecasting errors ($\times 100$) for the rolling window estimates, with a fitting period length of 66. Text in bold shows the model with lowest RMSE and MAE over the specific horizon. Number of forecasting periods, Δ, is equal to 96.

** Denotes rejection at 1% significance level.

* Denotes rejection at 5% significance level.

Table 12: One week forecasts using period specification δ

Panel A: In-sample estimation results

	EUR/SEK, One Week Forecast, $RV_{t+1,t+5}$					**USD/SEK**, One Week Forecast, $RV_{t+1,t+5}$				
	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV
c	0.008** (0.003)	0.009** (0.003)	0.007** (0.003)	0.009** (0.003)	0.007** (0.003)	0.009** (0.004)	0.003 (0.004)	0.004 (0.004)	0.014** (0.004)	0.009* (0.004)
RV_t	0.337** (0.048)	-	0.293** (0.050)	-	-	0.214** (0.048)	-	0.175** (0.050)	-	-
$RV_{t-4,t}$	0.205** (0.069)	-	0.166* (0.068)	-	-	0.309** (0.072)	-	0.278** (0.071)	-	-
$RV_{t-21,t}$	0.367** (0.060)	-	0.100 (0.085)	-	-	0.411** (0.062)	-	0.165* (0.081)	-	-
IV_t	-	0.982** (0.032)	0.400** (0.093)	-	0.369** (0.097)	-	1.020** (0.030)	0.369** (0.082)	-	0.279** (0.082)
C_t	-	-	-	0.414** (0.058)	0.387** (0.057)	-	-	-	0.237** (0.054)	0.203** (0.055)
$C_{t-4,t}$	-	-	-	0.119 (0.092)	0.063 (0.092)	-	-	-	0.405** (0.088)	0.379** (0.087)
$C_{t-21,t}$	-	-	-	0.390** (0.093)	0.125 (0.114)	-	-	-	0.323** (0.073)	0.144 (0.090)
J_t	-	-	-	-0.002 (0.060)	-0.030 (0.060)	-	-	-	-0.020 (0.044)	-0.023 (0.043)
$J_{t-4,t}$	-	-	-	0.193 (0.113)	0.208 (0.111)	-	-	-	-0.128 (0.088)	-0.123 (0.087)
$J_{t-21,t}$	-	-	-	-0.015 (0.162)	-0.010 (0.160)	-	-	-	0.010 (0.128)	0.069 (0.127)
$AdjR^2$	0.743	0.699	0.754	0.752	0.760	0.781	0.739	0.790	0.797	0.802

Panel B: Out-of-sample estimation results

	EUR/SEK, One Week Forecast, $RV_{t+1,t+5}$					**USD/SEK**, One Week Forecast, $RV_{t+1,t+5}$				
	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV
RMSE	2.031	1.983	1.987	1.989	**1.966**	2.564	2.318	2.437	2.312	**2.282**
MAE	1.210	1.236	**1.170**	1.224	1.188	1.666	**1.439**	1.571	1.517	1.467

Estimation results for HAR-RV equation (15), IV equation (17), HAR-IV equation (18), HAR-CJ equation (16) and HAR-CJ-IV equation (19). Panel A presents the in-sample regression estimates; the first column lists the independent variables while the remaining columns report the corresponding coefficients together with the standard errors in parentheses and the adjusted R^2 in the bottom row. Panel B reports the average RMSE and MAE one step-ahead forecasting errors ($\times$ 100) for the rolling window estimates, with fitting period length of 66. The number of forecasting periods, δ, is equal to 416.

** Denotes rejection at 1% significance level.

* Denotes rejection at 5% significance level.

Table 13: EUR/USD estimation results

Panel A: In-sample estimation results

	One Week Forecast, $RV_{t+1,t+5}$					One Month Forecast, $RV_{t+1,t+22}$				
	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV
c	0.010 (0.005)	0.014* (0.006)	0.010 (0.005)	0.015* (0.007)	0.013 (0.007)	0.016** (.006)	0.020** (.006)	0.016** (0.006)	0.026** (0.008)	0.024** (0.007)
RV_t	0.191 (0.104)	-	0.166 (0.108)	-	-	0.078 (0.118)	-	0.029 (0.121)	-	-
$RV_{t-4,t}$	0.222 (0.142)	-	0.221 (0.143)	-	-	0.031 (0.162)	-	0.029 (0.160)	-	-
$RV_{t-21,t}$	0.496** (0.117)	-	0.384** (0.169)	-	-	0.745** (0.133)	-	0.522** (0.190)	-	-
IV_t	-	0.817** (0.050)	0.134 (0.146)	-	0.101 (0.155)	-	0.774** (0.052)	0.265 (0.164)	-	0.154 (0.170)
C_t	-	-	-	0.195 (0.108)	0.173 (0.113)	-	-	-	0.116 (0.118)	0.085 (0.124)
$C_{t-4,t}$	-	-	-	0.158 (0.142)	0.159 (0.142)	-	-	-	-0.042 (0.156)	-0.040 (0.156)
$C_{t-21,t}$	-	-	-	0.555** (0.124)	0.464* (0.186)	-	-	-	0.814** (0.136)	0.675** (0.204)
J_t	-	-	-	-0.097 (0.097)	-0.092 (0.097)	-	-	-	-0.297** (0.106)	-0.290** (0.107)
$J_{t-4,t}$	-	-	-	0.396* (0.182)	0.405* (0.183)	-	-	-	0.313 (0.200)	0.328 (0.201)
$J_{t-21,t}$	-	-	-	-0.275 (0.267)	-0.251 (0.270)	-	-	-	-0.329 (0.293)	-0.293 (0.296)
$AdjR^2$	0.795	0.738	0.794	0.800	0.798	0.720	0.701	0.725	0.744	0.744

Panel B: Out-of-sample estimation results

	One Week Forecast, $RV_{t+1,t+5}$					One Month Forecast, $RV_{t+1,t+22}$				
	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV	HAR-RV	IV	HAR-RV-IV	HAR-CJ	HAR-CJ-IV
RMSE	**1.425**	1.859	1.512	1.532	1.598	**1.740**	1.931	1.817	1.834	1.909
MAE	**1.181**	1.400	1.205	1.258	1.273	**1.309**	1.487	1.360	1.319	1.410

Estimation results for the HAR-RV equation (15), IV equation (17), HAR-IV equation (18), HAR-CJ equation (16) and HAR-CJ-IV equation (19). Panel A presents the in-sample regression estimates; the first column lists the independent variables while the remaining columns report the corresponding coefficients together with their standard errors in parentheses and the adjusted R^2 in the bottom row. Panel B reports the average RMSE and MAE one step-ahead forecasting errors ($\times 100$) for the rolling window estimates, with a fitting period length of 66. Text in bold shows the model with lowest RMSE and MAE over the specific horizon. Number of forecasting periods, Δ, is equal to 97.
** Denotes rejection at 1% significance level.
* Denotes rejection at 5% significance level.

Figure 7: EUR/SEK: Comparison between original and robust time setting

(a) $RV_{t+1,t+5}$

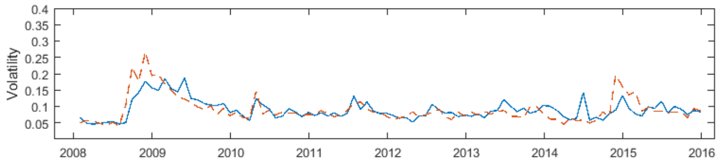

(b) $RV_{t+1,t+22}$

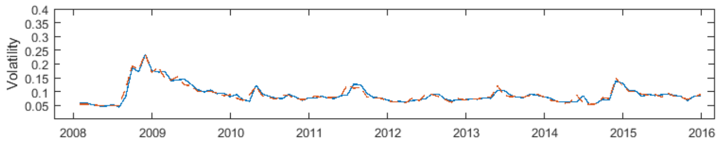

Comparison with the the original (solid) and robust (dashed) variable specifications displayed for $RV_{t+1,t+5}$ in (a) and $RV_{t+1,t+22}$ in (b).

Figure 8: USD/SEK: Comparison between original and robust time setting

(a) $RV_{t+1,t+5}$

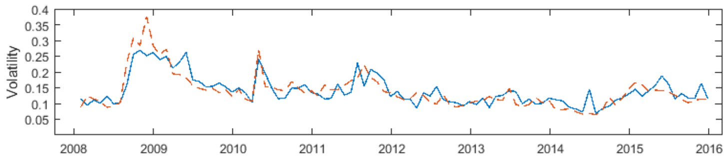

(b) $RV_{t+1,t+22}$

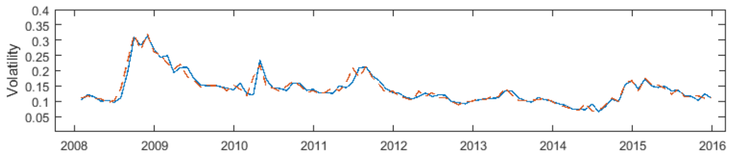

Comparison with the the original (solid) and robust (dashed) variable specifications displayed for $RV_{t+1,t+5}$ in (a) and $RV_{t+1,t+22}$ in (b).