



FX market volatility modelling: Can we use low-frequency data?

Štefan Lyócsa^{*,a,b}, Tomáš Plíhal^a, Tomáš Výrost^{a,c}

^a Institute of Financial Complex Systems, Masaryk University, Lipova 41a, Brno 602 00, Czech Republic

^b Faculty of Management, University of Presov, Konstantinova 16, Presov 080 01, Slovakia

^c University of Economics in Bratislava, Dolnozemska cesta 1, Bratislava 852 35, Slovakia

ARTICLE INFO

Keywords:

Volatility modelling
Foreign exchange markets
High-frequency data
HAR
Realized GARCH

ABSTRACT

High-frequency data tend to be costly, subject to microstructure noise, difficult to manage, and lead to high computational costs. Is it always worth the extra effort? We compare the forecasting accuracy of low- and high-frequency volatility models on the market of six major foreign exchange market (FX) pairs. Our results indicate that for short-forecast horizons, high-frequency models dominate their low-frequency counterparts, particularly in periods of increased volatility. With an increased forecast horizon, low-frequency volatility models become competitive, suggesting that if high-frequency data are not available, low-frequency data can be used to estimate and predict long-term volatility in FX markets.

1. Introduction

The outburst of the global financial markets in 2008, the European debt crisis, (geo)political uncertainties, oil-price wars in 2019 and 2020, and the outbreak of COVID-19 in 2020 have resulted in a surge in volatility in financial markets worldwide. For example, investors use volatility estimates for pricing financial derivatives. Fund managers might set specific risk levels that are, in turn, influenced by the predicted level of volatility. Risk levels are also targeted by banks to fulfill specific Basel criteria. Volatility might even be traded (using options or artificial indices linked to market volatility [Poon and Granger, 2003](#)). Times of extreme volatility also create pressure to rebalance portfolios, and the likelihood of contagion between markets also increases ([Kodres and Pritsker, 2002](#)). Market participants are thus interested in measuring, managing, and forecasting market volatility to determine the value of their investments and to prepare and communicate their planned market decisions.

The literature on volatility forecasting is rich and unfolds around available volatility estimators. Initially, volatility was calculated from low-frequency, daily data. The first generation of generalized autoregressive conditional heteroscedasticity (GARCH) models ([Bollerslev, 1986](#)) emerged in the 1990s and early 2000s and is represented by numerous variations using low-frequency data, e.g., EGARCH, GJR-GARCH, AP-ARCH, N-GARCH, NA-GARCH, I-GARCH, and FIGARCH (for an earlier review, see [Poon and Granger, 2003](#)). The GARCH class of models offers competitive forecasts and can capture many stylized facts about volatility, particularly the volatility clustering effect. With the greater availability of high-frequency data in the late 2000s, the research shifted toward high-frequency (intraday) volatility estimators and models.

The heterogeneous autoregressive (HAR) models of [Corsi \(2009\)](#) utilized high-frequency data and the realized volatility estimator of [Andersen and Bollerslev \(1998\)](#) and [Andersen et al. \(2001\)](#). The empirical evidence suggests that models of volatility based on high-frequency estimators provide superior forecasts to models based on low-frequency data (e.g., [Andersen et al., 2007](#); [Koopman](#)

* Corresponding author.

E-mail addresses: stefan.lyocsa@econ.muni.cz (Š. Lyócsa), tomas.plihal@econ.muni.cz (T. Plíhal), tomas.vyrost@euba.sk (T. Výrost).

et al., 2005; Corsi et al., 2010; Busch et al., 2011; Horpestad et al., 2019). Although the basic HAR model of Corsi (2009) is appealingly simple and appears to capture the short- and long-term dependency of the volatility process adequately (e.g., Andersen et al., 2007; Vortelinos, 2017), the literature has raised several issues related to the effect of microstructure noise.¹ Previously, Andersen et al. (2001) acknowledged that for realized volatility (high-frequency estimator of daily volatility) to be more efficient and unbiased, one needs high-quality data from actively traded assets. As a response, alternative estimators have emerged (e.g., Ait-Sahalia et al., 2005; Bandi and Russell, 2008; Barndorff-Nielsen et al., 2008; Andersen et al., 2011; Liu et al., 2015a).

The second generation of GARCH models bridges these two strands of the literature by relying on the latent volatility model (the GARCH concepts) while also using high-frequency data. The key ideas of the realized-GARCH model were presented by Hansen et al. (2012), and several alternative models emerged thereafter (e.g., Wu and Xie, 2019; Xie and Yu, 2019).

Despite the wide interest of academia, the existing literature provides evidence only that (i) volatility estimators based on high-frequency data are theoretically preferred (Andersen et al., 2001) and (ii) in the *day-ahead* predictive setting, models using high-frequency data provide superior performance (e.g., Andersen et al., 2007; Koopman et al., 2005; Corsi et al., 2010; Busch et al., 2011; Horpestad et al., 2019). Over longer horizons, averaging daily low-frequency volatility estimators across multiple days should reduce the effect of noise. Intuitively, intraday price fluctuations should not greatly contribute to month-ahead volatility forecasts. Therefore, with increasing forecast horizon, the difference in using high- or low-frequency volatility estimators should decrease, at which point low-frequency volatility models should tend to provide similarly accurate forecasts to high-frequency volatility models. Evidence on the relative (un)importance of low-frequency volatility models for *multiple-day-ahead* forecasts is *lacking*, which is intriguing, given that the heterogeneity of market participants has increased (with different needs and investment horizons, Woolbridge, 2019) and in many real-world scenarios, market participants are more interested in long-term forecasts, e.g., derivative traders. We fill this gap in the literature. In a recent study, Ma et al. (2018) showed that when low- and high-frequency volatility forecasts are combined appropriately, the accuracy increases for the Shanghai Stock Exchange Composite Index and S&P 500 index. Therefore, low-frequency data could provide additional information complementary to the available high-frequency data. Nevertheless, the study of Ma et al. (2018) is centered around day-ahead forecasts, where high-frequency volatility models should have the edge.

In this study, we present the results from a volatility forecasting modeling framework that compares the forecasting accuracy of several low- and high-frequency volatility models as a function of the forecast horizon. Our market of interest is represented by six major currency pairs.²

For some, the implications of our research could be substantial. If low-frequency volatility models provide competitive performance, one could argue that high-frequency data are not always worth the much higher costs. Daily foreign exchange data are freely available from various sources,³ but availability of high-frequency foreign exchange data depends on the policy of the given broker or bank, and data are not always free.⁴ Even if data are available⁵ for free, they are subject to various constraints, e.g., have limited licensing (e.g., can be used only for academic purposes) or are available only for short time periods or for a specific time frequency. Moreover, the use of high-frequency data raises other issues, most notably, working with high-frequency data requires appropriate cleaning and processing of the data. For example, the approximate sizes of the daily EUR/USD data from 2005 to 2019 is 120 kB, 5-second data is 350 MB, and tick-by-tick data is 15 GB. Processing daily data and estimating the models is overall much faster than processing and estimating models that use high-frequency data, where one needs to clean and prepare each line of the 15 GB of data.⁶ Therefore, the processing, data management, and computational intensity demands are much higher for high-frequency data and might not be worth the greater effort. Our results illustrate the dominance of high-frequency estimators for forecasting one-day-ahead volatility. Models that utilize high-frequency data or their combinations provide superior results. However, for longer forecast horizons, the combination of low-frequency volatility models provides forecasts statistically comparable to those of high-frequency volatility models and their combinations. Our results suggest that for most foreign exchange market (FX) pairs, low-frequency data represent a sufficient replacement for high-frequency data for forecast horizons of 5 or more days. Our study might therefore provide practitioners and policymakers with evidence supporting the use of high- or low-frequency volatility models in a particular setting.

¹ The basic specification of the HAR model has also been enhanced, e.g., by the inclusion of semivariances (Patton and Sheppard, 2015), the disentanglement of the realized volatility into continuous and jump components (e.g., Andersen et al., 2012), the introduction of the measurement error of the realized volatility into the HAR model as in Bollerslev et al. (2016), the inclusion of nontrading volatility components (Lyócsa and Molnár, 2017; Lyócsa and Todorova, 2020), and the use of hidden Markov chains (Luo et al., 2019).

² Equities and commodities are addressed in a separate study and show qualitatively similar results.

³ e.g., finance.yahoo.com, investing.com.

⁴ For example, the well-known provider of high-frequency data, Tick Data (www.tickdata.com), provides tick-by-tick quote data (bid and ask prices) that are already cleaned and processed. Moreover, these data are from more contributors (banks and other market participants). The dataset that we used in our paper would cost approximately 8 100 USD after all discounts (July 2020).

⁵ e.g., Oanda, dukascopy.

⁶ One needs to do this only once, but we want to stress that different types of skills and experience are also required to work with high-frequency data.

2. Methodology

2.1. Volatility estimators

2.1.1. High-frequency estimator

Given 5-minute intraday continuous returns $r_{t,j}$ for day $t = 1, 2, \dots, T$ and intraday period $j = 1, 2, \dots, N$, the usual realized variance estimator⁷ (e.g., Andersen and Bollerslev, 1998; Andersen et al., 2001) is defined as:

$$RV_t = \sum_{j=1}^N r_{t,j}^2 \quad (1)$$

Many alternative estimators of quadratic variation to address the inherent microstructure noise exist (e.g., Zhang et al., 2006; Jacod et al., 2009; Andersen et al., 2012). Our choice to use the 5-minute realized variance estimators is motivated by Liu et al. (2015b), who compared the empirical accuracy of several estimators across many assets⁸. They found that consistently outperforming the simple 5-minute realized variance is difficult.

2.1.2. Low-frequency estimators

As an alternative low-frequency estimator, we use range-based estimators that are more efficient than the usual daily squared return (e.g., Molnár, 2012). Motivated by Patton and Sheppard (2009), we increase the efficiency of the estimation process by combining three range-based estimators via a simple average. Specifically, given the natural logarithm of opening (O_t), high (H_t), low (L_t), and closing (C_t) prices on day t , the Parkinson (1980) estimator is:

$$PK_t = \frac{(H_t - L_t)^2}{4 \times \ln 2} \quad (2)$$

The Garman and Klass (1980) estimator is:

$$GK_t = 0.511(H_t - L_t)^2 - 0.019(C_t(H_t + L_t) - 2H_t L_t) - 0.383(C_t - O_t)^2 \quad (3)$$

Both estimators assume that the price follows driftless geometric Brownian motion. Allowing for arbitrary drift, Rogers and Satchell (1991) derived the following estimator:

$$RS_t = H_t(H_t - C_t) + L_t(L_t - C_t) \quad (4)$$

The range-based estimator used in our empirical setting is the average (following Patton and Sheppard, 2009) of the above three estimators:

$$RB_t = (PK_t + GK_t + RS_t)/3 \quad (5)$$

The motivation behind using the (naive) equally weighted average is based on the assumption that we have no prior information on which estimator might be more accurate for a given trading day.⁹ Should this simplified approach lead to competitive multiple-day-ahead volatility forecasts, it follows that a more sophisticated combination of low-frequency estimators might make the results even stronger.

2.2. Volatility models

In this section, we describe what we refer to as high- and low-frequency volatility models. As the name suggests, high-frequency volatility models utilize realized variance as the estimator of volatility, whereas low-frequency models use the range-based estimator. We use three classes of models: the heterogeneous autoregressive model (HAR) of Corsi (2009), the autoregressive fractionally integrated model (ARFIMA), and the realized generalized autoregressive conditional heteroscedasticity (realized-GARCH) of Hansen et al. (2012). These models were selected because they can use either high- or low-frequency volatility estimators in a straightforward manner. Moreover, all these models have been proven to be capable of replicating long memory and volatility clustering effects.

2.2.1. HAR class volatility models

In the past decade, the simple HAR model proposed by Corsi (2009) has gained popularity since it is easy to estimate and tends to perform better than competing first-generation GARCH models (Horpestad et al., 2019). Let $RV_{t+1,t+h}$ be the daily average realized variance calculated over the next h days. In this study, we are especially interested in the role of low-frequency estimators for

⁷ In the following text, we use the terms variance and volatility interchangeably.

⁸ Their comparison also included foreign exchange market futures.

⁹ The development and statistical verification of a method that continuously updates weights is left for further research. However, motivated by reviewer insights, we run our analysis and compare the results with low-frequency volatility models that use each of the three range-based estimators separately. A short discussion is presented in Section '4.3. Individual range-based low-frequency volatility forecasts'.

multiple-day-ahead volatility forecasts.

We employ 1-to-66 trading day-ahead forecasts. According to the recent Bank for International Settlements (BIS) survey, in 2019, 78% of the over-the-counter (OTC) foreign exchange derivatives had a maturity of less than one year.¹⁰ For the low-frequency volatility models to be useful for a wide array of participants, they should produce competitive forecasts up to a forecast horizon of one year or less. As our analysis shows that after a few weeks, the low-frequency volatility models tend to provide competitive forecasts across all FX pairs, we have used 66 trading days (three months) as a compromise between a few weeks and one year.

Our baseline HAR model is therefore specified as:

$$RV_{t+1,t+h} = \beta_1 + \beta_2 RV_t + \beta_3 RV_{t,t-4} + \beta_4 RV_{t,t-21} + \beta_5 RV_{t,t-65} + \epsilon_t \quad (6)$$

RV_t is the realized variance, and $RV_{t,t-4}$, $RV_{t,t-21}$, and $RV_{t,t-65}$ are average realized variances calculated over the past 5, 22 and 66 days, respectively.

The multiple-component volatility structure in (6) is motivated by the heterogeneous market hypothesis of Müller et al. (1997), according to which different market participants have various trading frequencies and time horizons, presumably because of different risk aversions, transaction costs, available information, and other constraints. Specifically, Corsi (2009, Section 2.2) argues for three volatility components: (i) short-term daily (1 day), (ii) medium-term weekly (5 days), and (iv) long-term corresponding to one (22 days) or more (e.g., 66 days) months. We follow the work of Corsi (2009) and use the daily, weekly and monthly volatility components. From the existing studies on foreign exchange market volatility forecasting (e.g., Bubák et al., 2011; Vortelinos, 2017), our specification differs only in that in addition to the one-month, we also incorporate the three-month volatility component, which is motivated by the fact that we are also predicting three-month (66-day) ahead volatility. The model is denoted *RV-HAR*, and the corresponding low-frequency, range-based version is denoted *RB-HAR*.

We consider two other popular versions of the HAR model that aim to model the asymmetric volatility observed in financial markets. Let NSV_t and PSV_t , respectively, denote the negative and positive semivariances of (e.g., Barndorff-Neilsen et al., 2010; Patton and Sheppard, 2015):

$$NSV_t = \sum_{j=2}^N r_{t,j}^2 I[r_{t,j} < 0], \quad PSV_t = \sum_{j=2}^N r_{t,j}^2 I[r_{t,j} > 0], \quad (7)$$

I represents an indicator function that returns one if the condition in square brackets holds and zero otherwise. The HAR model is then defined as:

$$RV_{t+1,t+h} = \beta_1 + \beta_2 NSV_t + \beta_3 PSV_t + \beta_4 RV_{t,t-4} + \beta_5 RV_{t,t-21} + \beta_6 RV_{t,t-65} + \epsilon_t \quad (8)$$

We use only one-day lags of NSV_t and PSV_t to mitigate the number of estimated parameters, which might deteriorate the forecasting performance in an out-of-sample context. Such simplified models were also considered by Patton and Sheppard (2015) and Bollerslev et al. (2016). This model is denoted *SV-RV-HAR*. As a low-frequency range-based counterpart, we use the following specification:

$$RB_{t+1,t+h} = \beta_1 + \beta_2 RB_t + \beta_3 RB_t \times I[R_t < 0] + \beta_4 RB_{t,t-4} + \beta_5 RB_{t,t-21} + \beta_6 RB_{t,t-65} + \epsilon_t \quad (9)$$

R_t is the daily return, and $\beta_3 RB_t \times I[R_t < 0]$ captures the asymmetric volatility response. The model is denoted *ARB-RB-HAR*.

The final two specifications are also motivated by the asymmetric volatility literature, namely, Corsi and Reno (2009) and Horpestad et al. (2019):

$$RV_{t+1,t+h} = \beta_1 + \beta_2 RV_t + \beta_3 |R_t| + \beta_4 |R_t| \times I[R_t < 0] + \beta_5 RV_{t,t-4} + \beta_6 RV_{t,t-21} + \beta_7 RV_{t,t-65} + \epsilon_t \quad (10)$$

The coefficient β_4 captures the asymmetric effect, and β_3 controls for the size effect. As argued by Horpestad et al. (2019), if absolute returns are correlated with variance (which is likely), one should also include $|R_t|$ in the equation. This model is denoted *L-RV-HAR*, and the range-based counterpart is denoted *L-RB-HAR*. All HAR models are estimated via weighted least squares, where the weights are reciprocal values of the dependent variable (see Clements and Preve, 2019 for a discussion of estimating HAR models).

2.2.2. ARFIMA-GARCH-class volatility model

We next use an ARFIMA-GARCH model, for which the mean equation models variance:

$$RV_t = \alpha + u_t \quad (11)$$

$$\left(1 - \sum_{i=1}^p \phi_i L^i\right) (1-L)^d u_t = \left(1 + \sum_{j=1}^q \theta_j L^j\right) \epsilon_t \quad (12)$$

$$\epsilon_t = v_t \eta_t, \text{ where } \eta_t \sim iid(0, 1) \quad (13)$$

¹⁰ BIS OTC derivative statistics are available at <https://stats.bis.org/statx/srs/table/d9?f=pdf>

where d is the differencing parameter (e.g., Granger and Joyeux, 1980), v_t is the time-varying volatility¹¹ and η_t is an iid variable following a flexible distribution (Johnson, 1949b; 1949a). The variance equation is the exponential GARCH model of Nelson (1991):

$$\ln v_t^2 = \omega + \alpha z_{t-1} + \gamma(|z_{t-1}| - E|z_{t-1}|) + \beta \ln v_{t-1}^2 \quad (14)$$

The sign and the size effects are captured by α and γ , and z_t is the standardized innovation. The high-frequency volatility model employs the realized variance and is denoted *RV-ARFIMA-GARCH*, and the range-based estimator is denoted *RB-ARFIMA-GARCH*.

2.2.3. Measurement equation GARCH volatility model

Finally, due to their popularity and the development of more sophisticated second-generation GARCH models, we use the realized-GARCH model of Hansen et al. (2012), which can be adjusted to work with high- or low-frequency volatility estimators. The mean equation models daily returns:

$$R_t = \kappa + u_t \quad (15)$$

$$\left(1 - \sum_{i=1}^p \phi_i L^i\right) (1-L)^d u_t = \left(1 + \sum_{j=1}^q \theta_j L^j\right) \varepsilon_t \quad (16)$$

$$\varepsilon_t = v_t \eta_t, \text{ where } \eta_t \sim iid(0, 1) \quad (17)$$

The variance and the measurement equations are:

$$\ln v_t^2 = \omega + \alpha \ln RV_{t-1} + \beta \ln v_{t-1}^2 \quad (18)$$

$$\ln RV_t = \xi + \delta \ln v_t^2 + \lambda_1 z_t + \lambda_2 (z_t^2 - 1) + w_t, \quad w_t \sim N(0, \theta) \quad (19)$$

Originally, Hansen et al. (2012) used realized variance, in which case we denote the model as *realized-GARCH*. If the range-based estimator is used instead, the model is called *range-GARCH*.

2.3. Forecasting procedure

The forecasting procedure uses a rolling-window framework. The algorithm is as follows:

1. Select observations from $t = 1, 2, \dots, T_e$.
2. Estimate volatility models.
3. Using estimated parameters and observations, predict volatility at $T_e + 1$. For *HAR* models, multiple-day-ahead forecasts are predicted directly, while for *ARFIMA-GARCH* and *real-GARCH* models, multiple-day-ahead forecasts are calculated recursively.
4. Shift the estimation window by using observations $t = 2, 3, \dots, T_e + 1$ and repeat steps 2 to 4 until the end of the sample.

The estimation window size is set to $T_e = 1000$.

2.4. Combinations of forecasts: discounted forecast errors

We draw on the ideas of Bates and Granger (1969) and use simple combination techniques to mitigate model uncertainty (Timmermann, 2006). Forecasts are combined across all high-frequency volatility models, all low-frequency volatility models, and all ten high- and low-frequency volatility models. To combine forecasts, we use weighted averages, where the weights are given by the discounted forecast error.

Let F_t^m and F_t denote the forecasts from model m and the corresponding proxy, the realized variance RV_t . Our first combination is a simple average across all forecasts:

$$C_H^{\text{Ave}} = M^{-1} \sum_{m=1}^M F_t^m \quad (20)$$

Here, the subscript H means that we averaged across high-frequency models. For low-frequency models, we use the subscript L , and for a combination across both classes of forecasts, we use HL . The loss (to be defined in the next section) is $L_t(F_t^m, F_t)$ and for simplicity is denoted as L_t . We use the discounted forecast error to weight each loss value such that recent losses have higher weight than losses in the past, and we calculate the average loss over a time period of T (out-of-sample) observations:

¹¹ In this case, it is the time-varying volatility of variance.

$$\bar{L} = \sum_t^T \frac{\delta^{T-t+1}}{\sum_k^T \delta^{T-k+1}} L_t \quad (21)$$

where δ is the weighting parameter. With $\delta = 1$, all losses have equal weights. The lower δ is, the higher the relative weight of the most recent losses. We choose $\delta = 0.975$ and observe almost no qualitative change in results for $\delta = 0.950$ or $\delta = 0.900$. The weighted losses are calculated from the most recent 200 predictions, which we refer to as the size of the *calibration sample*. Thus, the first combination forecast is available for the 1201st observation of the initial sample (Estimation window + calibration sample + 1). Our second combination is formed as the weighted trimmed mean:

$$C_H^{Trim} = \sum_{(m)=2}^{M-1} \bar{L}^{(m),*} F_t^{(m)} \quad (22)$$

Here, $F_t^{(m)}$ represents the ordered forecasts, i.e., the lowest and the highest are excluded, and $\bar{L}^{(m),*}$ are the corresponding losses that are rescaled to sum to one. The final combination is a weighted average across the three best performing models and is denoted C_H^{Top} .

2.5. Volatility forecast evaluation

As noted in the previous section, our proxy is the realized variance, RV_b , which in the subsequent equations is denoted F_t . This approach clearly places low-frequency models at a disadvantage, but we argue that this is the only meaningful way to test whether low-frequency models can achieve comparable performance to that of high-frequency models.

We evaluate the forecasts of our model specification using two statistical loss functions and the model confidence set (MCS). According to Patton (2011), the mean square error (MSE) and quasi-likelihood (QLIKE) loss functions provide a consistent ranking of forecasts, even if the proxy of the underlying latent volatility is measured with noise.

$$L_t^{QLIKE} = \frac{F_t}{F_t^m} - \ln \frac{F_t}{F_t^m} - 1 \quad (23)$$

$$L_t^{MSE} = (F_t^m - F_t)^2 \quad (24)$$

As the QLIKE loss function is less sensitive to extreme values and penalizes underestimation of volatility more strongly, we use it to present our key results.¹²

Statistical evaluation is conducted based on the MCS proposed by Hansen et al. (2011). This algorithm is suitable when models are nested, when a benchmark model is not specified, and when multiple models are evaluated, i.e., it controls for data-snooping bias. The MCS algorithm finds the 'superior set of models', which represents models with the same predictive ability at the selected confidence level.

3. Data

We study the market with the largest turnover in the world (approximately 6.59 trillion per day in April 2019, on average), the foreign exchange market, specifically, the six most liquid currency pairs, namely, AUD/USD, EUR/USD, GBP/USD, USD/CAD, USD/CHF, and USD/JPY. Our selected currency pairs represent approximately 30% of the entire market turnover.¹³ Our sample covers fifteen years of data, covering the 5th of May 2005 to the 24th of September 2019.

We collect data from OANDA using a 5-minute calendar sampling scheme over a 24-hour trading window that starts at 22:00 UTC (end of the New York session). Due to low liquidity, weekends are removed from the analysis to avoid estimation bias, as is standard in the literature (e.g., Dacorogna et al., 2001; Andersen et al., 2007; Aloud et al., 2013; Gau and Wu, 2017).

The descriptive statistics for our daily volatility measures and returns are presented in Table 1. We note several interesting differences between high- and low-frequency variance estimates. First, the distribution of the low-frequency variance estimates shows a higher spread of values, which we would expect from a noisier estimate. Specifically, the low-frequency variance estimate has an approximately 30% larger standard deviation with more skew and higher kurtosis. Second, on average, the low-frequency estimator is slightly smaller than its high-frequency counterpart.¹⁴ Third, the persistence of the high-frequency estimators is higher and shows longer memory. This characteristic might prove to be useful in HAR models, which specifically exploit this persistence. Fourth, the correlation between daily high- and low-frequency variance estimators is 0.90 (AUD/USD), 0.86 (EUR/USD), 0.96 (GBP/USD), 0.83 (USD/CAD), 0.88 (USD/CHF), and 0.89 (USD/JPY). Given that these correlations are based on daily estimates, we consider these values to be sufficiently high to warrant a meaningful volatility comparison, which is the subject of this study.

¹² Qualitatively, we do not obtain different conclusions when interpreting the results using MSE. The results are available as supplementary electronic material.

¹³ See https://www.bis.org/statistics/rpfx19_fx.pdf

¹⁴ Subtracting the two estimators and regressing against the constant shows that this difference is statistically significant for AUD/USD and USD/CAD at the 0.01 level, for GBP/USD and USD/JPY at the 0.05 level, and not significant for EUR/USD and USD/CHF.

Table 1

Descriptive statistics of the variance estimators and daily returns of FX rates.

FX pair	Mean	SD	Skew.	Kurt.	$\rho(1)$	$\rho(5)$	$\rho(22)$	$\rho(66)$
Realized variance (annualized)								
AUD/USD	175.716	359.817	11.465	207.707	0.777	0.623	0.429	0.200
EUR/USD	91.211	100.924	5.075	46.771	0.696	0.553	0.437	0.332
GBP/USD	99.987	226.420	27.277	1015.974	0.292	0.215	0.159	0.131
USD/CAD	98.820	111.833	4.747	39.784	0.779	0.689	0.574	0.372
USD/CHF	110.333	219.861	24.140	807.654	0.402	0.189	0.122	0.077
USD/JPY	109.210	187.484	11.675	218.589	0.453	0.244	0.164	0.102
Range-based variance (annualized)								
AUD/USD	164.898	391.774	13.081	253.819	0.607	0.463	0.314	0.151
EUR/USD	90.188	126.985	6.376	66.144	0.435	0.339	0.274	0.198
GBP/USD	96.890	289.675	28.858	1066.861	0.188	0.111	0.102	0.091
USD/CAD	90.024	138.368	10.080	192.082	0.481	0.387	0.284	0.208
USD/CHF	107.714	346.444	41.415	2139.290	0.190	0.078	0.040	0.028
USD/JPY	105.514	222.498	13.503	276.415	0.291	0.144	0.093	0.053

Note: $\rho(\cdot)$ is the value of the auto-correlation coefficient at the given lag. The SD is the standard deviation. The correlation between high- and low-frequency variance estimators is 0.90, 0.86, 0.96, 0.83, 0.88, and 0.89 for AUD/USD, EUR/USD, GBP/USD, USD/CAD, USD/CHF, USD/JPY.

4. Results

4.1. Comparing high- and low-frequency volatility forecasts

For illustration purposes, Figs. 1 and 2 plot the daily realized variance for the six FX pairs and corresponding day-ahead forecasts from ARFIMA-GARCH models, which tend to produce the most accurate day-ahead forecasts for both high- and low-frequency volatility models. The forecasts tend to follow realized variances but are unable to replicate sudden spikes in volatility, a phenomenon also visible in other forecasting studies.

Our key results visualized in Fig. 3 (QLIKE losses) show the lowest forecast error that can be achieved within either high-frequency (black line) or low-frequency (red line) forecasting models as a function of the forecast horizon.¹⁵ A dot on the line highlights that the given model belongs to the set of superior models. These figures illustrate our key observation that for short-forecast horizons, high-frequency models tend to be superior in terms of mean forecast errors, but this advantage disappears for longer forecast horizons, where the forecasts errors are statistically indistinguishable.

Specific numerical results are presented in Tables 2–5. These tables contain the average of the QLIKE loss function for 1-, 5-, 22-, and 66-day-ahead forecast horizons. The values in bold and with the dagger symbol represent the models that belong to the MCS, i.e., the predictive abilities of the models in bold are considered to be equally good. For example, the best models for forecasting one-day volatility for EUR/USD (second column in Table 2) are C_H^{Ave} and C_H^{Trim} , which combine the results from high-frequency models (Panel C).

For weekly volatility forecasts (see Table 3), the combinations of low-frequency volatility models in Panel D provide competitive forecasts to those of the high-frequency models for three FX pairs, GBP/USD, USD/CHF and USD/JPY. For monthly and quarterly volatility forecasts (Table 4) we find suitable low-frequency alternatives to high-frequency models for all FX pairs except USD/CAD. The difficulty in finding competitive low-frequency volatility models for USD/CAD is not surprising given that for USD/CAD, the low-frequency estimators had the lowest correlation with their high-frequency counterparts.

Combining individual forecasts proves to be beneficial for all forecast horizons. However, for monthly and quarterly forecast horizons, several individual low-frequency forecasting models show good performance that is statistically indistinguishable from that of the high-frequency models (Panel D in Tables 2–5). The results for the MSE loss function lead to similar conclusions.¹⁶ An exception is that the low-frequency volatility models are also competitive for the USD/CAD FX pair for monthly and quarterly volatility forecasts.

With respect to individual models, we find that for short-term forecasts, the HAR-based models tend to underperform the ARFIMA-GARCH and realized-GARCH models. As the forecasting horizon increases, the accuracy of the HAR models also increases, which demonstrates how ignoring model uncertainty can influence the conclusions from such studies. Following the earlier work of Ma et al. (2018), we also study whether combining high- and low-frequency volatility models can further improve the accuracy of forecasting models, but we were unable to confirm this hypothesis (see the results in Panel E, Tables 2–5).

In summary, as we increase the forecasting horizon, distinguishing between high- and low-frequency volatility models becomes more difficult. This finding can be generalized for all FX pairs and for both loss functions. Moreover, while combining high- and low-frequency volatility forecasts has not proven to be particularly useful, combining forecasts only from high-frequency volatility models or only from low-frequency volatility models is a good strategy regardless of the forecasting horizon and volatility estimator employed.

¹⁵ To facilitate a better comparison, both figures present forecast errors from only the best performing high- and low-frequency volatility models, i.e., from the top performers. Note that the top-performing models are not visible in these figures and that they might change with the forecast horizon.

¹⁶ see the Electronic supplementary material for the corresponding figures and tables.

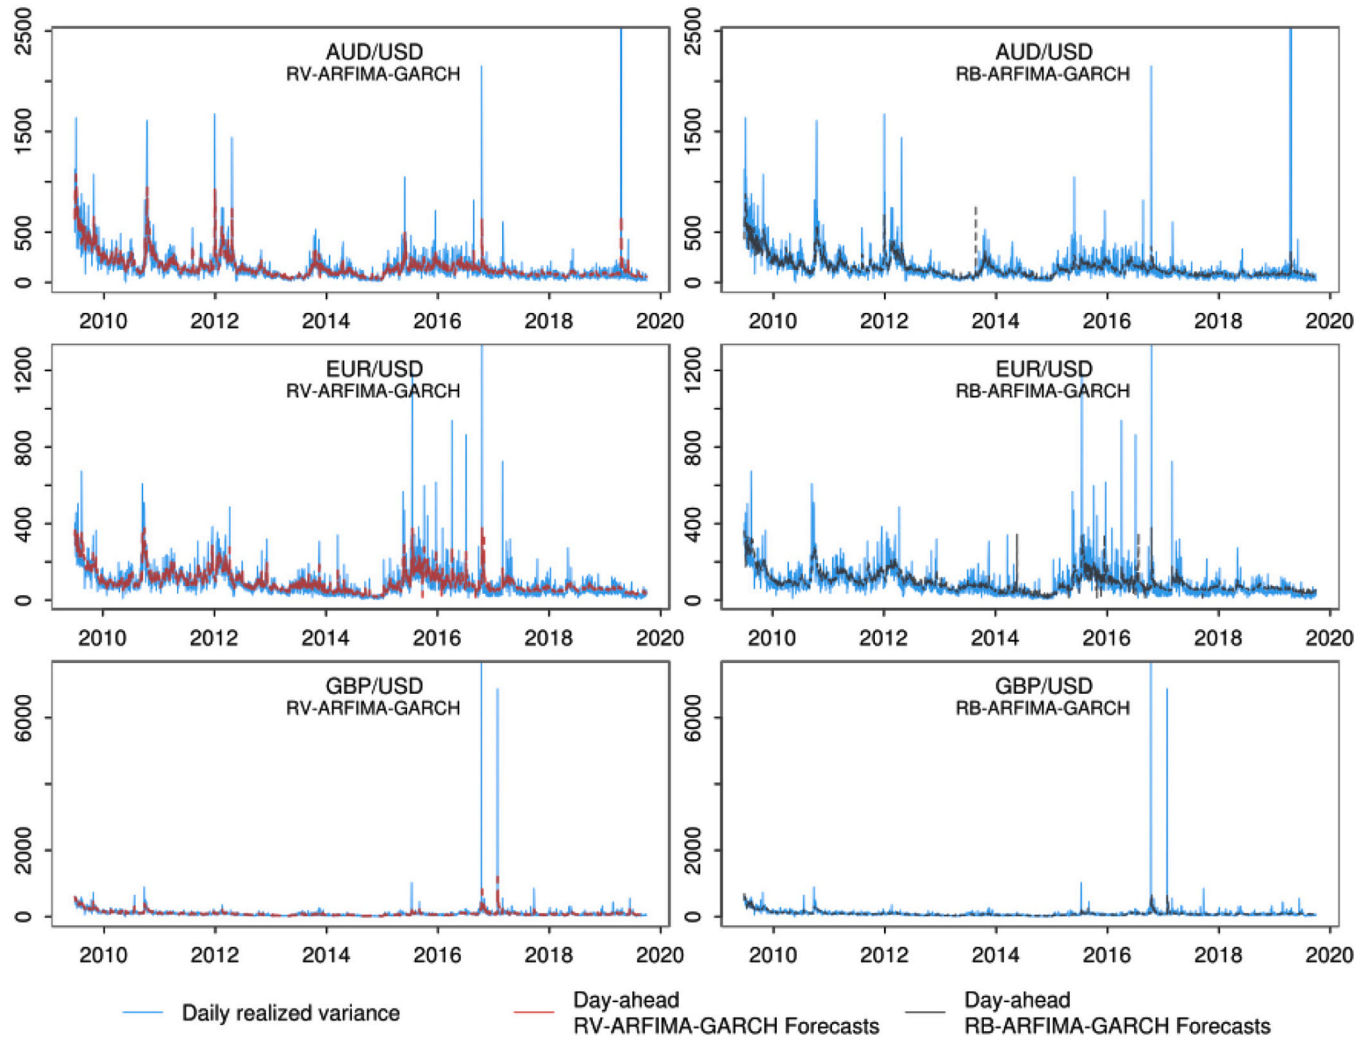


Fig. 1. Comparison of the RV-ARFIMA-GARCH and RB-ARFIMA-GARCH models with realized variance.

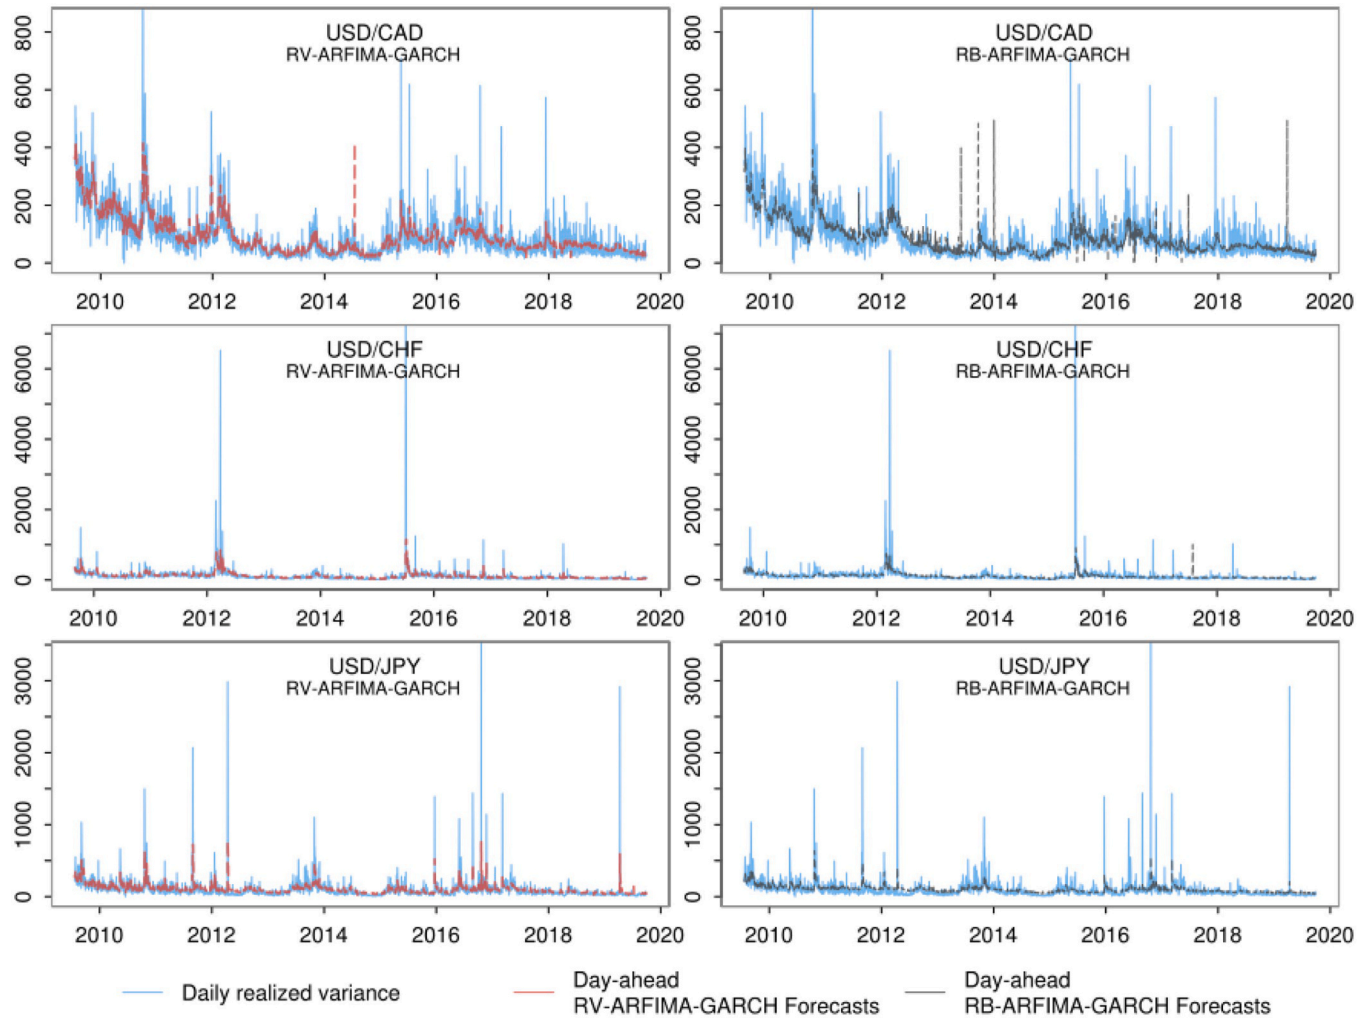


Fig. 2. Comparison of the RV-ARFIMA-GARCH and RB-ARFIMA-GARCH models with realized variance.




Fig. 3. High- and low-frequency volatility forecast QLIKE loss functions for different forecasting horizons.

Table 2

Average QLIKE loss function for 1-day-ahead forecasts.

FX pair	AUD/USD	EUR/USD	GBP/USD	USD/CAD	USD/CHF	USD/JPY
Panel A: Individual high-frequency volatility model forecasts						
RV-HAR	0.337	0.060	0.121	7.691	0.095	0.258
SV-RV-HAR	1.958	0.060	0.112	2.260	0.081†	0.244
L-RV-HAR	0.318	0.060	0.125	1.283	0.091	0.318
RV-ARFIMA-GARCH	0.065	0.073	0.059†	0.075†	0.088	0.142
realized-GARCH	0.067	0.081	0.072	0.090	0.085	0.084†
Panel B: Individual low-frequency volatility model forecasts						
RB-HAR	0.257	0.225	0.922	0.248	0.287	0.323
ARB-RB-HAR	0.254	0.236	0.824	0.253	0.278	0.302
L-RB-HAR	0.261	0.222	1.521	0.244	0.293	0.305
RB-ARFIMA-GARCH	0.092	0.148	0.096	0.105	0.138	0.198
range-GARCH	0.082	0.115	0.098	0.105	0.127	0.109
Panel C: Combining high-frequency volatility forecasts						
C_H^{Ave}	0.077	0.052†	0.071	0.123	0.072†	0.094
C_H^{Trim}	0.058†	0.052†	0.068	0.082	0.072†	0.084†
C_H^{Top}	0.058†	0.054	0.059†	0.083	0.068†	0.082†
Panel D: Combining low-frequency volatility forecasts						
C_L^{Ave}	0.112	0.115	0.187	0.135	0.142	0.129
C_L^{Trim}	0.085	0.100	0.092	0.112	0.113	0.118
C_L^{Top}	0.086	0.105	0.109	0.116	0.119	0.110
Panel E: Combining high- and low-frequency volatility forecasts						
C_{HL}^{Ave}	0.084	0.072	0.103	0.108	0.092	0.104
C_{HL}^{Trim}	0.067	0.062	0.080	0.093	0.084	0.094
C_{HL}^{Top}	0.062	0.055	0.062†	0.074†	0.068†	0.089

Notes: The values in bold and with † symbol denote model confidence set for given currency pair. In other words, we can not reject the hypothesis that these models have the same predictive performance at the level of $\alpha = 0.15$. All models and forecast combinations are described in Section 2.

Table 3

Average QLIKE loss function for 5-day-ahead forecasts.

FX pair	AUD/USD	EUR/USD	GBP/USD	USD/CAD	USD/CHF	USD/JPY
Panel A: Individual high-frequency volatility model forecasts						
RV-HAR	0.066†	0.065†	0.098†	0.043†	0.101	0.180
SV-RV-HAR	0.066†	0.065†	0.098†	0.044†	0.096†	0.177
L-RV-HAR	0.066†	0.064†	0.094†	0.042†	0.100†	0.171
RV-ARFIMA-GARCH	0.065†	0.068†	0.076†	0.044†	0.096†	0.167
realized-GARCH	0.075†	0.078	0.081†	0.067	0.098†	0.122†
Panel B: Individual low-frequency volatility model forecasts						
RB-HAR	0.094	0.087	0.139	0.078	0.127	0.200
ARB-RB-HAR	0.094	0.086	0.137	0.078	0.124	0.191
L-RB-HAR	0.094	0.086	0.136	0.077	0.128	0.191
RB-ARFIMA-GARCH	0.072	0.088	0.100	0.062	0.116	0.190
range-GARCH	0.074	0.083	0.105†	0.064	0.125	0.123†
Panel C: Combining high-frequency volatility forecasts						
C_H^{Ave}	0.061†	0.059†	0.080†	0.041†	0.087†	0.139†
C_H^{Trim}	0.062†	0.060†	0.083†	0.042†	0.090†	0.141†
C_H^{Top}	0.061†	0.061†	0.081†	0.041†	0.087†	0.142†
Panel D: Combining low-frequency volatility forecasts						
C_L^{Ave}	0.074	0.072	0.105	0.061	0.099†	0.142†
C_L^{Trim}	0.069	0.072	0.097†	0.059	0.090†	0.138†
C_L^{Top}	0.070	0.073	0.102†	0.062	0.093†	0.146†
Panel E: Combining high- and low-frequency volatility forecasts						
C_{HL}^{Ave}	0.064†	0.063†	0.089†	0.047	0.090†	0.140†
C_{HL}^{Trim}	0.064†	0.063†	0.092†	0.046	0.092†	0.143
C_{HL}^{Top}	0.061†	0.062†	0.083†	0.042†	0.090†	0.137†

Notes: The values in bold and with † symbol denote model confidence set for given currency pair. In other words, we can not reject the hypothesis that these models have the same predictive performance at the level of $\alpha = 0.15$. All models and forecast combinations are described in Section 2.

Table 4

Average QLIKE loss function for 22-day-ahead forecasts.

FX pair	AUD/USD	EUR/USD	GBP/USD	USD/CAD	USD/CHF	USD/JPY
Panel A: Individual high-frequency volatility model forecasts						
RV-HAR	0.070†	0.063†	0.104†	0.043†	0.129†	0.156†
SV-RV-HAR	0.069†	0.063†	0.102†	0.043†	0.129†	0.156†
L-RV-HAR	0.069†	0.062†	0.102†	0.043†	0.130†	0.155†
RV-ARFIMA-GARCH	0.071†	0.072†	0.105†	0.048†	0.140†	0.168†
realized-GARCH	0.088†	0.086	0.112†	0.071	0.150†	0.134†
Panel B: Individual low-frequency volatility model forecasts						
RB-HAR	0.084	0.075	0.124	0.060	0.140†	0.161†
ARB-RB-HAR	0.083	0.075	0.126	0.060	0.137†	0.158†
L-RB-HAR	0.083	0.073†	0.120	0.059	0.141†	0.159†
RB-ARFIMA-GARCH	0.072†	0.090	0.117†	0.063	0.144†	0.170†
range-GARCH	0.085	0.080	0.127†	0.069	0.160	0.132†
Panel C: Combining high-frequency volatility forecasts						
C_H^{Ave}	0.067†	0.060†	0.099†	0.042†	0.125†	0.139†
C_H^{Trim}	0.067†	0.061†	0.101†	0.043†	0.126†	0.143†
C_H^{Top}	0.069†	0.061†	0.102†	0.042†	0.126†	0.147†
Panel D: Combining low-frequency volatility forecasts						
C_L^{Ave}	0.073†	0.067†	0.112†	0.054	0.124†	0.138†
C_L^{Trim}	0.072†	0.072	0.112†	0.057	0.125†	0.144†
C_L^{Top}	0.076†	0.072†	0.117	0.057	0.123†	0.147†
Panel E: Combining high- and low-frequency volatility forecasts						
C_{HL}^{Ave}	0.068†	0.062†	0.103†	0.045†	0.123†	0.138†
C_{HL}^{Trim}	0.068†	0.063†	0.105†	0.045†	0.126†	0.143†
C_{HL}^{Top}	0.069†	0.062†	0.104†	0.047†	0.131†	0.146†

Notes: The values in bold and with † symbol denote model confidence set for given currency pair. In other words, we can not reject the hypothesis that these models have the same predictive performance at the level of $\alpha = 0.15$. All models and forecast combinations are described in Section 2.

Table 5

Average QLIKE loss function for 66-day-ahead forecasts.

FX pair	AUD/USD	EUR/USD	GBP/USD	USD/CAD	USD/CHF	USD/JPY
Panel A: Individual high-frequency volatility model forecasts						
RV-HAR	0.078†	0.070†	0.092†	0.061†	0.151†	0.132†
SV-RV-HAR	0.078†	0.070†	0.091†	0.061†	0.150†	0.131†
L-RV-HAR	0.078†	0.068†	0.091†	0.061†	0.150†	0.132†
RV-ARFIMA-GARCH	0.077†	0.083†	0.103†	0.066†	0.165	0.138†
realized-GARCH	0.102	0.092	0.110†	0.082	0.153†	0.126†
Panel B: Individual low-frequency volatility model forecasts						
RB-HAR	0.087†	0.076†	0.100†	0.075	0.154†	0.127†
ARB-RB-HAR	0.087†	0.076†	0.101†	0.075	0.154†	0.125†
L-RB-HAR	0.087†	0.074†	0.101†	0.075	0.154†	0.127†
RB-ARFIMA-GARCH	0.077†	0.089	0.109†	0.080	0.157†	0.134†
range-GARCH	0.099	0.079†	0.112†	0.084	0.163†	0.123†
Panel C: Combining high-frequency volatility forecasts						
C_H^{Ave}	0.077†	0.071†	0.091†	0.061†	0.148†	0.122†
C_H^{Trim}	0.078†	0.069†	0.093†	0.061†	0.149†	0.124†
C_H^{Top}	0.078†	0.067†	0.093†	0.061†	0.148†	0.125†
Panel D: Combining low-frequency volatility forecasts						
C_L^{Ave}	0.080†	0.072†	0.096†	0.071	0.142†	0.116†
C_L^{Trim}	0.081†	0.069†	0.096†	0.075	0.136†	0.116†
C_L^{Top}	0.084†	0.069†	0.097†	0.074	0.149†	0.121†
Panel E: Combining high- and low-frequency volatility forecasts						
C_{HL}^{Ave}	0.077†	0.072†	0.092†	0.063†	0.143†	0.119†
C_{HL}^{Trim}	0.078†	0.068†	0.093†	0.064†	0.148†	0.121†
C_{HL}^{Top}	0.079†	0.062†	0.092†	0.067†	0.149†	0.121†

Notes: The values in bold and with † symbol denote model confidence set for given currency pair. In other words, we can not reject the hypothesis that these models have the same predictive performance at the level of $\alpha = 0.15$. All models and forecast combinations are described in Section 2.

4.2. Dissecting good and bad volatility performance

In this section, we explore differences in forecasting performance. First, note that the results in Table 1 show that the low-frequency variance estimates are somewhat lower. Thus, the forecasts from low-frequency volatility models might be (downward) biased. Therefore, we conduct the (Mincer and Zarnowitz, 1969) test. Specifically, we estimate a regression model of the form:

$$F_t = \beta_0 + \beta_1 \hat{F}_t + u_t \quad (25)$$

and conduct a joint test¹⁷ of $\beta_0 = 0$, $\beta_1 = 1$. The rejection of the test suggests biased forecasts.

For short-forecast horizons, the ARFIMA-GARCH and realized-GARCH models do not produce biased forecasts, while the HAR models do. As the forecast horizon increases, the HAR models become less biased. The combination forecasts appear to be biased most of the time for most forecast horizons, which is interesting because the theoretical superiority of combination forecasts is derived from uncorrelated and unbiased forecasts being combined (averaged). The results from the Mincer and Zarnowitz (1969) tests show that combination forecasts appear to be biased most of the time for most forecast horizons. As bias is deemed a negative feature of forecasts, it is surprising to observe that combination forecasts actually lead to more accurate forecasts (see Panels C and D in Tables 2–5). These results are similar for both high- and low-frequency volatility models, although for shorter forecast horizons, low-frequency models tend to be more biased, which suggests why high-frequency models tend to be more accurate for shorter forecast horizons. Intuitively, in the MSE framework, such a result is possible if the increased bias is compensated by lower variance of forecasts. This results also suggests that resolving the bias might lead to further improvements in the forecasting accuracy.¹⁸

Finally, we study how the difference between the accuracy of high- and low-volatility forecasting models changes over time and in high-/low-volatility periods. For this analysis, we compare the accuracy of forecasts generated from the C^{Trim} combination forecasts, which generally lead to the most competitive forecasts for both high- and low-frequency volatility models. The results are based on the QLIKE loss function. We run the following regression:

$$L_t^{C^{Trim}} - L_t^{C^{Trim}} = LD_t^{C^{Trim}} = \beta_0 + \beta_1 RV_{t-1} + \beta_2 t + u_t \quad (26)$$

The larger the loss differential $LD_t^{C^{Trim}}$ is, the more accurate the high-frequency forecast C_H^{Trim} . If β_0 is positive, then the high-frequency model tends to be systematically more accurate. If β_1 is positive, then the accuracy of the high-frequency models is higher during periods of higher volatility, while a nonzero β_2 coefficient suggests that the accuracy changes systematically over time.

The estimated coefficients reported in Table 6 show that over time, high-frequency models tend to produce more precise forecasts for the AUD/USD, EUR/USD, GBP/USD, and USD/CAN FX pairs, while the accuracy is also increased during more volatile periods. The opposite is true for USD/JPY, and the results are nonsignificant for USD/CHF. These results suggest that more accurate forecasting models could be designed with a conditional combination that would exploit the level of market volatility.

4.3. Individual range-based low-frequency volatility forecasts

Up to now, for our low-frequency volatility models, we have assumed that we do not have any ex ante information about which of the range-based estimators leads to more accurate volatility forecasts. Here, we discuss the results from low-frequency volatility models estimated separately for the Garman and Klass (1980), Parkinson (1980) and Rogers and Satchell (1991) estimators. Detailed tabulated results are available upon request.

Our general observation does not change. Increasing the forecast horizon leads to more competitive forecasts from low-frequency volatility models regardless of the range-based estimator employed. Among the individual range-based estimators, the Garman and Klass (1980) estimator leads to lower forecast errors compared to the forecast errors generated from volatility models based on the equally weighted average of the range-based estimators.

However, this does not mean that one should blindly prefer the Garman and Klass (1980) estimator, as there are two caveats. First, using only one range-based estimators has occasionally led to very inaccurate forecasts, which could successfully be avoided by using the average of the three range-based estimators. For example, in a day-ahead setting for the GBP/USD and USD/CAD pairs, the forecast errors from the RB-HAR models with the Garman and Klass (1980) estimator lead to 1.231 and 7.325 average QLIKE losses, in contrast to 0.922 and 0.248 when employing the average. Second, the results considerably differ across FX pairs and forecast horizons. For example, in 1- and 5-day-ahead settings, the use of the Garman and Klass (1980) estimator leads to worse forecasts for the USD/JPY pair but to more accurate forecasts for 22- and 66-day-ahead settings. For EUR/USD, the trend is reversed, with Garman and Klass (1980) being more accurate for 1- and 5-day-ahead settings and worse for 22- and 66-day-ahead forecasts.

These examples suggest that in many practical scenarios, using the average across estimators should be preferred to using individual estimators.

¹⁷ The significance of the test is based on a variance-covariance matrix estimated with a quadratic spectral weighting scheme and the automatic bandwidth selection of (Newey and West, 1994).

¹⁸ We leave this option for future research and do not explore it further here.

Table 6

Conditions under which high-frequency models tend to outperform low-frequency models.

	AUD/USD		EUR/USD		GBP/USD		USD/CAD		USD/CHF		USD/JPY	
	Coef.	Sig.	Coef.	Sig.	Coef.	Sig.	Coef.	Sig.	Coef.	Sig.	Coef.	Sig.
<i>Panel A: day-ahead forecasts</i>												
Constant	163.44	**	478.3	***	121.1		−60.6		217.37	***	653.83	***
RV_{t-1}	0.13		0.71		−0.2	*	0.82		0.97	***	−0.32	
Trend	0.12		−0.02		0.09	*	0.25	***	0.1		−0.2	***
R^2	0.11%		0.08%		0.15%		0.87%		1.00%		0.53%	
<i>Panel B: five-day-ahead forecasts</i>												
Constant	−188.86		−262.14		−999.88	*	−174.02	**	756.46		728.71	*
RV_{t-1}	1.82	**	3.28	**	12.73	**	3.63	***	−4.19		−4.91	
Trend	0.02		0.1		0.1		0.06	*	−0.28		−0.28	**
R^2	4.83%		5.92%		31.05%		7.70%		4.95%		6.33%	
<i>Panel C: twenty-two-day-ahead forecasts</i>												
Constant	−345.06	***	−161.76		−726.59	***	−344.13	***	1291.06		582.29	***
RV_{t-1}	2.82	***	1.96	**	9.29	***	4.16	***	−7.76		−4	***
Trend	0.04		0.08	*	0.07		0.11	***	−0.46		−0.18	***
R^2	12.32%		2.08%		33.03%		10.03%		13.55%		12.06%	
<i>Panel D: sixty-six-day-ahead forecasts</i>												
Constant	−226.72	***	237.04		−145.54		−228.85	**	950.82		507.08	***
RV_{t-1}	1.94	***	−0.64		1.58		3.51	***	−4.66		−4.2	***
Trend	0.01		0.01		0.06	**	0.08	**	−0.26		−0.17	***
R^2	6.96%		0.04%		1.71%		6.13%		2.31%		10.92%	

Note: The results correspond to the modelling of the loss differential between C_H^{Trim} and C_L^{Trim} forecasting models by the means of lagged realized variance and trend variable. All coefficients are multiplied by 10^4 . Significances are based on the variance-covariance matrix estimated using a quadratic spectral weighting scheme and Newey and West automatic bandwidth selection. */**/** correspond to 10%, 5% and 1% significance levels.

5. Conclusion

As many subjects interact with the FX market, predicting the market's uncertainty is crucial for improved risk management. While high-frequency data lead to superior volatility estimates, the acquisition, data management and computational costs associated with such data cannot be covered by all market participants. Moreover, low-frequency data are publicly available and are much easier to work with. This leads to our question from the title 'Can we use low-frequency data?'.

In this paper, we compare the forecasting performance of several volatility models that use either low- or high-frequency volatility estimates or both. On the basis of a sample of six major currency pairs, our results suggest that for short-forecast horizons (from 1 to 5 days), high-frequency models dominate their low-frequency counterparts. As the forecast horizon increases, the advantage of the high-frequency models disappears, and low- and high-frequency forecasts become statistically comparable. The answer to the question proposed in the title is 'if high-frequency data are not available, then low-frequency data can be used to estimate and predict long-term market volatility'.

Moreover, regardless of whether one relies on high- or low-frequency volatility models, one should utilize combination forecasts. The Mincer and Zarnowitz (1969) tests further suggest that at least part of the inaccuracy of low-frequency volatility forecasts is due to bias. Finally, we find that high-frequency models tend to be more superior during periods of increased volatility.

These results have implications for researchers and investors alike, as they demonstrate that low-frequency volatility models can provide competitive performance to that of high-frequency models under some circumstances. Our study notes that high-frequency data might not always be worth the much higher acquisition, data management and processing costs, especially if the forecast horizon of interest is sufficiently long.

CRediT authorship contribution statement

Štefan Lyócsa: Conceptualization, Methodology, Software, Data curation, Formal analysis, Writing - original draft, Writing - review & editing, Visualization, Funding acquisition, Project administration. **Tomáš Plíhal:** Conceptualization, Methodology, Data curation, Investigation, Writing - original draft, Writing - review & editing. **Tomáš Výrost:** Methodology, Software, Validation, Writing - review & editing, Formal analysis.

Acknowledgement

This research was supported by the Czech Science Foundation (GACR), nr. 18-05829S. We would like to thank Oleg Deev for his research assistance and useful comments.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at [10.1016/j.frl.2020.101776](https://doi.org/10.1016/j.frl.2020.101776)

References

- Ait-Sahalia, Y., Mykland, P.A., Zhang, L., 2005. How often to sample a continuous-time process in the presence of market microstructure noise. *Rev. Financ. Stud.* 18 (2), 351–416.
- Aloud, M., Fasli, M., Tsang, E., Dupuis, R., Olsen, R., 2013. Stylized facts of trading activity in the high frequency FX market: an empirical study. *J. Finance Invest. Anal.* 2 (4), 145–183.
- Andersen, T., Bollerslev, T., Diebold, F., Ebens, H., 2001. The distribution of realized stock return volatility. *J. Financ. Econ.* 61 (1), 43–76.
- Andersen, T.G., Bollerslev, T., 1998. Answering the skeptics: yes, standard volatility models do provide accurate forecasts. *Int. Econ. Rev.* 39, 885–905.
- Andersen, T.G., Bollerslev, T., Diebold, F.X., 2007. Roughing it up: including jump components in the measurement, modeling, and forecasting of return volatility. *Rev. Econ. Stat.* 89 (4), 701–720.
- Andersen, T.G., Bollerslev, T., Meddahi, N., 2011. Realized volatility forecasting and market microstructure noise. *J. Econom.* 160 (1), 220–234.
- Andersen, T.G., Bollerslev, T., 2012. Jump-robust volatility estimation using nearest neighbor truncation. *J. Econom.* 169 (1), 75–93.
- Bandi, F.M., Russell, J.R., 2008. Microstructure noise, realized variance, and optimal sampling. *Rev. Econ. Stud.* 75 (2), 339–369.
- Barndorff-Nielsen, O.E., Kinnebrock, S., Shephard, N., 2010. In: Bollerslev, T., Russell, J., Watson, M. (Eds.). *Measuring Downside Risk: Realised Semivariance*. Oxford University Press, pp. 117–136.
- Barndorff-Nielsen, O.E., Hansen, P.R., Lunde, A., Shephard, N., 2008. Designing realized kernels to measure the ex post variation of equity prices in the presence of noise. *Econometrica* 76 (6), 1481–1536.
- Bates, J.M., Granger, C.W., 1969. The combination of forecasts. *J. Oper. Res. Soc.* 20 (4), 451–468.
- Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. *J. Econom.* 31 (3), 307–327.
- Bollerslev, T., Patton, A.J., Quaedvlieg, R., 2016. Exploiting the errors: a simple approach for improved volatility forecasting. *J. Econom.* 192 (1), 1–18.
- Bubák, V., Kočenda, E., Žikeš, F., 2011. Volatility transmission in emerging European foreign exchange markets. *J. Bank. Finance* 35 (11), 2829–2841.
- Busch, T., Christensen, B.J., Nielsen, M.O., 2011. The role of implied volatility in forecasting future realized volatility and jumps in foreign exchange, stock, and bond markets. *J. Econom.* 160 (1), 48–57.
- Clements, A., Preve, D., 2019. A practical guide to harnessing the har volatility model. Available at SSRN 3369484.
- Corsi, F., 2009. A simple approximate long-memory model of realized volatility. *J. Financ. Econom.* 7 (2), 174–196.
- Corsi, F., Pirino, D., Reno, R., 2010. Threshold bipower variation and the impact of jumps on volatility forecasting. *J. Econom.* 159 (2), 276–288.
- Corsi, F., Reno, R., 2009. Har volatility modelling with heterogeneous leverage and jumps. Available at SSRN 1316953.
- Dacorogna, M., Müller, U., Pictet, O., de Vries, C., 2001. Extremal forex returns in extremely large data sets. *Extremes* 4 (2), 105.
- Garman, M.B., Klass, M.J., 1980. On the estimation of security price volatilities from historical data. *J. Bus.* 53 (1), 67.
- Gau, Y.-F., Wu, Z.-X., 2017. Macroeconomic announcements and price discovery in the foreign exchange market. *J. Int. Money Finance* 79, 232–254.
- Granger, C.W., Joyeux, R., 1980. An introduction to long-memory time series models and fractional differencing. *J. Time Ser. Anal.* 1 (1), 15–29.
- Hansen, P.R., Huang, Z., Shek, H.H., 2012. Realized GARCH: a joint model for returns and realized measures of volatility. *J. Appl. Econom.* 27 (6), 877–906.
- Hansen, P.R., Lunde, A., Nason, J.M., 2011. The model confidence set. *Econometrica* 79 (2), 453–497.
- Horpestad, J.B., Lyócsa, Š., Molnár, P., Olsen, T.B., 2019. Asymmetric volatility in equity markets around the world. *North Am. J. Econ. Finance* 48, 540–554.
- Jacod, J., Li, Y., Mykland, P.A., Podolskij, M., Vetter, M., 2009. Microstructure noise in the continuous case: the pre-averaging approach. *Stoch. Process. Appl.* 119 (7), 2249–2276.
- Johnson, N.L., 1949. Bivariate distributions based on simple translation systems. *Biometrika* 36 (3/4), 297–304.
- Johnson, N.L., 1949. Systems of frequency curves generated by methods of translation. *Biometrika* 36 (1/2), 149–176.
- Kodres, L.E., Pritsker, M., 2002. A rational expectations model of financial contagion. *J. Finance* 57 (2), 769–799.
- Koopman, S.J., Jungbacker, B., Hol, E., 2005. Forecasting daily variability of the S&P 100 stock index using historical, realised and implied volatility measurements. *J. Empir. Finance* 12 (3), 445–475.
- Liu, L.Y., Patton, A.J., Sheppard, K., 2015. Does anything beat 5-minute RV? A comparison of realized measures across multiple asset classes. *J. Econom.* 187 (1), 293–311.
- Liu, L.Y., Patton, A.J., Sheppard, K., 2015. Does anything beat 5-minute RV? A comparison of realized measures across multiple asset classes. *J. Econom.* 187 (1), 293–311.
- Luo, J., Klein, T., Ji, Q., Hou, C., 2019. Forecasting realized volatility of agricultural commodity futures with infinite hidden markov har models. *Int. J. Forecast.* https://www.sciencedirect.com/science/article/pii/S0169207019302316?casa_token=rtU5iTpqkAAAAA:heWVDHu5SSZFXD_BQs_c-Gb1SeLU_k5yaoqK_bFaEXWOD_Ocjrqn_fBQsTfKqelncJ0gTjAM.
- Lyócsa, Š., Molnár, P., 2017. The effect of non-trading days on volatility forecasts in equity markets. *Finance Res. Lett.* 23, 39–49.
- Lyócsa, Š., Todorova, N., 2020. Trading and non-trading period realized market volatility: does it matter for forecasting the volatility of us stocks? *Int. J. Forecast.* 36 (2), 628–645.
- Ma, F., Li, Y., Liu, L., Zhang, Y., 2018. Are low-frequency data really uninformative? Aforecasting combination perspective. *North Am. J. Econ. Finance* 44, 92–108. <https://doi.org/10.1016/j.najef.2017.11.006>.
- Mincer, J.A., Zarnowitz, V., 1969. The evaluation of economic forecasts. *Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Performance*. NBER, pp. 3–46.
- Molnár, P., 2012. Properties of range-based volatility estimators. *Int. Rev. Financ. Anal.* 23, 20–29.
- Müller, U.A., Dacorogna, M.M., Davé, R.D., Olsen, R.B., Pictet, O.V., Von Weizsäcker, J.E., 1997. Volatilities of different time resolutions-analyzing the dynamics of market components. *J. Empir. Finance* 4 (2-3), 213–239.
- Nelson, D.B., 1991. Conditional heteroskedasticity in asset returns: a new approach. *Econometrica* 347–370.
- Newey, W.K., West, K.D., 1994. Automatic lag selection in covariance matrix estimation. *Rev. Econ. Stud.* 61 (4), 631–653.
- Parkinson, M., 1980. The extreme value method for estimating the variance of the rate of return. *J. Bus.* 61–65.
- Patton, A.J., 2011. Volatility forecast comparison using imperfect volatility proxies. *J. Econom.* 160 (Realized Volatility), 246–256.
- Patton, A.J., Sheppard, K., 2009. Optimal combinations of realised volatility estimators. *Int. J. Forecast.* 25 (2), 218–238.
- Patton, A.J., Sheppard, K., 2015. Good volatility, bad volatility: Signed jumps and the persistence of volatility. *Rev. Econ. Stat.* 97 (3), 683–697.
- Poon, S.-H., Granger, C.W.J., 2003. Forecasting volatility in financial markets: a review. *J. Econ. Lit.* 41 (2), 478–539.
- Rogers, L.C.G., Satchell, S.E., 1991. Estimating variance from high, low and closing prices. *Ann. Appl. Probab.* 504–512.
- Timmermann, A., 2006. Forecast combinations. *Handb. Econ. Forecast.* 1, 135–196.
- Vortelinos, D.I., 2017. Forecasting realized volatility: har against principal components combining, neural networks and GARCH. *Res. Int. Bus. Finance* 39, 824–839.
- Wooldridge, P.D., December 2019. Fx and OTC derivatives markets through the lens of the triennial survey. *BIS Q. Rev.*

- Wu, X., Xie, H., 2019. A realized EGARCH-midas model with higher moments. *Finance Res. Lett.* 101392.
- Xie, H., Yu, C., 2019. Realized GARCH models: Simpler is better. *Finance Res. Lett.*
- Zhang, L., et al., 2006. Efficient estimation of stochastic volatility using noisy observations: A multi-scale approach. *Bernoulli* 12 (6), 1019–1043.