The Information Content of Short-Term Options*

October 19, 2018

Abstract

We exploit a recent development in the S&P 500 index option market to compute the weekly implied variance index. We show that the weekly implied variance is a strong predictor of the weekly realized variance. In an encompassing regression test, it crowds out the information content of the monthly implied variance. Further tests reveal that the weekly implied variance outperforms not only the monthly implied variance but also recently proposed time series models of realized variance. This result holds both in- and out-of-sample and the forecast accuracy gains are significant.

JEL classification: G11, G12

Keywords: Implied Variance, Predictability, Realized Variance, Weekly Options

*We are grateful to Torben Andersen, Fabian Hollstein, Marcel Prokopczuk, Rogier Quaedvlieg, Viktor Todorov and participants at the American Economic Association (2018), the British Accounting and Finance Association (2018), the Western Economic Association International Conference (2018), the Econometrics and Financial Data Science Workshop (2017), the EFiC Conference in Banking and Finance (2017), the International Conference on Computational and Financial Econometrics (2017), the Financial Econometrics Conference (2018), Leibniz University Hannover, the University of Glasgow, the University of Kent and the University of Stirling for helpful comments and discussions.

1 Introduction

Several studies, e.g. Jiang and Tian (2005), Carr and Wu (2009) and Busch et al. (2011), document that the variance implied by monthly options is a strong predictor of the monthly realized variance.[1] While this fact is generally well-accepted, we know relatively little about the forecasting power of short-term implied variance for short-term realized variance. Although it is tempting to speculate that the findings will hold for the shorter horizon, there are a couple of reasons to suggest otherwise. First, the literature routinely discards short-term options on the grounds that they are illiquid and noisy. Consequently, their information content is expected to be limited. Second, for the sample period of most existing studies, e.g. Jiang and Tian (2005) and Busch et al. (2011), there were very few short-term options available in the market, making it challenging to obtain a long enough time series. This limitation is important because the statistical tests may lack power in a short sample. As a work around this issue, some studies, e.g. Blair et al. (2001), Pong et al. (2004) and Kourtis et al. (2016), have resorted to using the monthly implied variance series to predict short-term realized variance. Alas, this approach introduces a mismatch between the maturity of the options and the forecasting horizon. It is unclear how big of an issue this disconnect may be.

We exploit an important innovation in the S&P 500 index option market, namely the launch of weekly option contracts (weeklies) by the Chicago Board of Options Exchange (CBOE) in 2005, to study the predictability of weekly realized variance. Essentially, these option contracts enable market participants to better manage their short-term risk, e.g. the weekly realized variance. Andersen et al. (2017) document that weeklies account for nearly 50 % of the total trading volume in the S&P 500 index options in 2015, indicating

[1] See also the survey by Poon and Granger (2003).

that these options are liquid. We use the Bakshi et al. (2003) estimator to compute the weekly option implied variance index and analyze its predictive power for the weekly realized variance estimated using 5-minute S&P 500 index data. In a regression of the daily time series of the weekly realized variance on a constant and the lagged weekly implied variance, we obtain a statistically significant slope estimate and a high predictive power ($Adj\ R^2 = 64.9\%$). In an effort to understand the channel through which this predictability arises, we decompose the realized variance into its continuous and jump components (Barndorff-Nielsen, 2002). We find that the weekly implied variance strongly predicts both the continuous ($Adj\ R^2 = 62.0\%$) and jump ($Adj\ R^2 = 47.3\%$) variations. Next, we evaluate the incremental information content of the weekly implied variance relative to that of the monthly implied variance. To this end, we formulate and estimate an encompassing model that includes the implied variance of both maturities, i.e. weekly and monthly. The regression results suggest that the weekly implied variance crowds out the monthly implied variance series. Furthermore, the forecasting performance of the weekly implied variance is significantly better than that of time series models of the heterogeneous autoregressive (HAR) realized variance family. This set of results holds both in- and out-of-sample.

We run a battery of tests to assess the robustness of our findings. The findings are unaffected when using a sampling frequency of 1-minute to compute the realized variance. Our conclusion is also robust to the method of interpolation used to compute the weekly and monthly implied variance series. Furthermore, the key findings are not driven by the choice of the implied variance estimator. Indeed, we obtain qualitatively similar results when using the Britten-Jones and Neuberger (2000) implied variance estimator.

Our research relates to the stream of the literature that uses the implied variance of a given maturity to predict the realized variance of a shorter horizon. Blair et al. (2001), Pong et al. (2004) and Kourtis et al. (2016) are interesting studies along those lines. We confirm their finding that the monthly implied variance predicts the weekly realized variance. However, we find that there are significant gains in forecasting accuracy once the maturity of the option contract and the forecasting horizon are aligned. In fact, our encompassing regression estimates suggest that the weekly implied variance crowds out the monthly implied variance. To the best of our knowledge, we are the first to document this result at the short horizon.

Our study adds to the broader literature on the predictability of realized variance. Corsi (2009) proposes the HAR model and documents its superior performance relative to the random walk model. Andersen et al. (2007) decompose the historical variance terms of the HAR model into continuous and jump components. Patton and Sheppard (2015) propose an extension that separately uses positive and negative semivariances. Bollerslev et al. (2016) extend the HAR model to account for heteroskedastic measurement errors in the realized variance series. Empirically, we show that the weekly implied variance provides significantly more accurate forecasts of short-term risk than the HAR model and its aforementioned extensions.

The remainder of this paper proceeds as follows. Section *2* introduces the methodology and the dataset. Section *3* discusses the performance of the weekly implied variance relative to that of the monthly implied variance series. Section *4* extends the analysis to various time series models of the HAR family. Section *5* presents the robustness checks. Finally, Section *6* concludes.

2 Data and Methodology

This section introduces the methodology used to construct the main variables used in our analysis. It then presents the dataset.

2.1 Methodology

2.1.a Realized Variance

Our paper focuses on the predictability of next week's realized variance. It is useful to start with the definition of the intraday return:

$$r_{j,k} = \log\left(\frac{S_{j,k}}{S_{j,k-1}}\right) \tag{1}$$

where $r_{j,k}$ denotes the intraday return at the end of the k^{th} intraday interval of day j. $S_{j,k}$ and $S_{j,k-1}$ are the asset prices at the end of the k^{th} and $(k-1)^{th}$ intraday interval of the trading day j, respectively.

We compute the (annualized) weekly realized variance as follows:

$$RV_{t+7}^{w} = 52 \times \sum_{j=0}^{N_{t+7}^{w}-1} \sum_{k=1}^{m} r_{t+7-j,k}^{2} \tag{2}$$

where RV_{t+7}^{w} is the (annualized) weekly realized variance for the week ending on day $t+7$. The number 52 indicates that the realized variance estimate is annualized. There are m returns observed each trading day. The case where $k = 1$ corresponds to the overnight return.[2] N_{t+7}^{w} indicates the number of trading days during the week ending on day $t+7$.

[2]It is standard in the literature to account for the overnight returns. See Bollerslev et al. (2009), Drechsler and Yaron (2011) and Bekaert and Hoerova (2014), among others.

2.1.b Implied Variance

The literature often uses the Britten-Jones and Neuberger (2000) estimator to proxy for the risk-neutral expectation of the total variation of returns. However, this estimator captures the risk-neutral expectation of the continuous variation which is equal to the total variation of returns only if the underlying return process does not jump. This result arises because the total variation is the sum of the continuous variation and the jump variation. In the presence of jumps, the Britten-Jones and Neuberger (2000) estimator is a biased estimator of the risk-neutral expectation of the future total variation with the magnitude of the bias increasing with the contribution of jumps to the total variation of returns (Du and Kapadia, 2013).

Andersen et al. (2007) and Lee and Mykland (2008), among others, use non-parametric statistical tests to show that the S&P 500 index jumps.[3] Du and Kapadia (2013) conduct an extensive simulation exercise to assess the impact of jumps on the implied variance series and recommend the Bakshi et al. (2003) estimator as a jump-robust estimator of implied variance. Heeding on their recommendation, we use the Bakshi et al. (2003) formula to compute the implied variance:[4]

$$IV_t^{\tau} = \frac{360}{\tau}\left[e^{\mathrm{rf}_t \frac{\tau}{360}}\mathrm{QUAD}_t - \left(e^{\mathrm{rf}_t \frac{\tau}{360}} - 1 - \frac{e^{\mathrm{rf}_t \frac{\tau}{360}}}{2}\mathrm{QUAD}_t\right)^2\right] \tag{3}$$

where

$$\mathrm{QUAD}_t = \int_0^{S_t} \frac{2\left(1+\ln\left[\frac{S_t}{K}\right]\right)}{K^2} P_t(\tau, K)dK + \int_{S_t}^{+\infty} \frac{2\left(1-\ln\left[\frac{K}{S_t}\right]\right)}{K^2} C_t(\tau, K)dK \tag{4}$$

[3] The documented jumps are not spuriously induced by the fact that the index is not directly tradable. Prokopczuk and Wese Simen (2016) show that the liquid S&P 500 E-Mini futures contract, which is tradable, jumps.

[4] To make our analysis more comparable to most studies, we also consider the Britten-Jones and Neuberger (2000) implied variance. Section *5.3* presents these results. Note that these findings should be interpreted cautiously, keeping in mind that this specific estimator is not robust to jumps in the underlying return process.

IV_t^τ is the (annualized) implied variance of time to maturity τ (expressed in days) observed at time t. Throughout this paper, we use the expressions "weekly" and "monthly" to denote the case where $\tau = 7$ and $\tau = 30$, respectively. $\mathrm{rf_t}$ is the τ-day (annualized) discount rate on day t. S_t is the underlying price at time t. $P_t(\tau, K)$ and $C_t(\tau, K)$ denote the price at time t of the European put and call options of time to maturity τ and strike price K, respectively. Note that the formula in Equation (4) involves only out-of-the-money (OTM) options. For each option maturity available on that day, we compute the Black and Scholes (1973) implied volatility of all OTM options. We then average the OTM implied volatility estimates of the same maturity. Equipped with this average implied volatility, denoted σ, we define the variables $K_{t,L}$ and $K_{t,U}$ as follows:

$$K_{t,L} = S_t e^{-8\sigma_t} \quad (5)$$

$$K_{t,U} = S_t e^{8\sigma_t} \quad (6)$$

where σ_t is the average implied volatility at time t of all OTM options of the same maturity.

Similar to Carr and Wu (2009), we linearly interpolate the implied volatilities for 2,000 equally spaced strike prices between $K_{t,L}$ and $K_{t,U}$ defined in Equations (5) and (6). In practice, the strike prices traded in the market do not completely span the interval starting at $K_{t,L}$ and ending at $K_{t,U}$, raising the question of extrapolation. We follow Jiang and Tian (2005) and Carr and Wu (2009), among others, and perform the nearest neighbourhood extrapolation. To be precise, for strike prices greater (lower) than $K_{t,L}$ ($K_{t,U}$) but lower (higher) than the lowest (highest) strike available in the market, we use the implied volatility associated with the lowest (highest) strike available in the market.

Next, we map the grid of 2,000 implied volatilities into Black and Scholes (1973) OTM option prices. Finally, we use the trapezoidal rule to numerically evaluate the integrals in Equation (4) and compute the implied variance as in Equation (3).[5]

We repeat the steps above for all maturities observed on that day, thus yielding the term-structure of implied variance. From this term-structure, we linearly interpolate the implied variance of weekly (IV^w) and monthly (IV^m) horizons. We emphasize that we only interpolate between maturities and never extrapolate since this could introduce spurious spikes in the constant maturity implied variance series.[6]

2.2 Data

We obtain high-frequency data on the S&P 500 index from Thomson Reuters Tick History (TRTH) to build the realized variance series. Our interest in high-frequency data, as opposed to daily data, is motivated by the studies of Andersen and Bollerslev (1998), Barndorff-Nielsen and Shephard (2002) and Andersen et al. (2003), who recommend the use of intraday data to efficiently measure realized variance. The dataset spans the period extending from January 1996 to August 2015.[7] It contains bid and ask quotes pertaining to regular business hours, i.e. from 08:30 AM to 3:00 PM (Chicago Time). Similar to Bollerslev et al. (2009) and Bollerslev and Todorov (2011), among others, we use a 5-minute sampling frequency.[8] At the end of each 5-minute interval, we use the most recent

[5] Note that by using options with strike prices ranging from $K_{t,L}$ to $K_{t,U}$, we essentially truncate the integrals in Equation (4). This choice is consistent with earlier work, e.g. Carr and Wu (2009). In a simulation setting, Jiang and Tian (2005) show that the truncation error is negligible if the truncation points are more than two standard deviations from the current underlying price.

[6] As a robustness check, we do not linearly interpolate the constant maturity contracts but instead simply use the option contracts with maturity closest to the target maturity. On average, the "real" maturities of the options are 6 and 27 calendar days for the weekly and monthly horizons, respectively. Our main conclusions are unchanged, leading us to conclude that the method of interpolation plays a minimal role in our results. See Section *5.2* for further details.

[7] The sample period is dictated by the dataset available from TRTH at the time we started this project.

[8] As a robustness check, we consider a higher sampling frequency of 1-minute and obtain qualitatively similar findings. See Section *5.1* for further details.

mid-quote price to proxy for the closing price of that interval.

We also obtain end-of-day S&P 500 index options data for the same period from IvyDB OptionMetrics. These options trade on the CBOE and are of the European type. For each trading day and option contract, the database contains information about the bid and ask prices, the open interest, the strike price and the expiration date. The dataset includes weekly and standard option contracts, among others. Generally, weekly options expire on the Friday of each week, except the third Friday of each month when the standard options expire.[9,10] Figure 3 of Andersen et al. (2017) reveals a rapid growth in the trading volume of the weekly contracts from less than 5 % of the total S&P 500 index option volume during the first few years of trading to 50 % towards the end of our sample.[11]

Although the OptionMetrics dataset spans the same time period as the TRTH dataset, we face two sample limitations that require us to start our analysis in March 2008. The first limitation is forced upon us since weekly options on the S&P 500 index were launched in October 2005. As a result, we can only analyze the period beginning from that point onwards. The second limitation is driven by the way OptionMetrics reports the closing option and underlying prices. Prior to the 5^{th} of March 2008, the OptionMetrics dataset records derivatives prices at the market closing (3:15 PM) whereas the underlying spot price is recorded at 3:00 PM, introducing a bias in studies that rely on synchronous

[9] At the time we started the project, the term-structure of weekly options includes up to 12 maturities. For further information about weekly options, we refer the interested reader to the following webpage: *http://www.cboe.com/micro/weeklys/introduction.aspx*.

[10] For an up-to-date list of weekly option contracts on offer, we direct the reader to the following link: *http://www.cboe.com/micro/weeklys/availableweeklys.aspx*.

[11] Absent official data on the identity/profile of market participants who trade weekly options, it is difficult to definitely ascertain their trading motives. Andersen et al. (2017) do not find a significant change in the trading activity of these contracts around important macroeconomic news announcements. This finding suggests that speculation does not seem to be the main driver of trading activity. It leaves open the possibility that the increased trading activity in weeklies is primarily driven by a desire to improve short-term risk management.

observations of the derivatives and spot prices. To remediate to this issue, we focus on the sample period extending from the 5^{th} of March 2008 to the end of August 2015.[12]

We process the option data as follows. We discard observations with missing or zero prices. We implement this filter using bid and ask prices separately. In doing so, we aim to tackle the concern that our dataset includes contracts that are not actively quoted. As is standard in the literature (Carr and Wu, 2009), we compute the mid-quote price of the option, which we refer to as the option price. Next, we remove all option observations that are in-the-money. We take this step because the computation of the implied variance only involves OTM option prices (see Equation (4)). Furthermore, we download the discount rates from OptionMetrics. These discount rates are based upon the London Interbank Offered Rates (LIBOR) and the Eurodollar futures. For each trading day and option contract, we linearly interpolate the discount rate of the same time to maturity as the option contract. We then match the discount rates with the panel of options data. We also match the time series of the daily S&P 500 index prices and that of the dividend yield, both obtained from OptionMetrics, with the panel of options data.

Our analysis involves daily observations of all key variables. Table 1 presents the summary statistics of the implied variance series. The weekly and monthly implied variance have average (annualized) values of 5.960% and 5.977%, respectively.

[12]One may argue that we should start our sample period at a later date, e.g. in 2011 as in Andersen et al. (2017), to allow the trading activity in weekly options to pick up. We also considered this alternative starting date and reached qualitatively similar results. If one holds the view that our current sample period includes illiquid weeklies, then this low trading activity should work against the predictive power of the weekly options. Viewed in this way, the gains in forecasting performance documented in our paper represent a "worst case" scenario.

3 Weekly vs. Monthly Implied Variance

This section starts by establishing the in-sample predictive power of the weekly implied variance for next week's realized variance and compares it to that of the monthly implied variance. We then explore the channels through which the predictability result arises.

3.1 Is Weekly Implied Variance Informative About Realized Variance?

3.1.a Univariate Evidence

We begin by evaluating the information content of implied variance for the week ahead realized variance. To this end, we estimate the following Mincer and Zarnowitz (1969) regression:

$$RV^{w}_{t+7} = \alpha + \beta IV^{x}_{t} + \epsilon_{t+7} \tag{7}$$

where α is the intercept. β denotes the slope parameter. IV^{x}_{t} is the implied variance on day t of time to maturity x, where x can be the weekly (w) or monthly (m) maturity. ϵ_{t+7} is the residual of the regression at $t+7$. If the implied variance is informative about the future weekly realized variance, then the slope parameter will be significantly different from 0.

Starting with the monthly implied variance, Table 2 reports a positive (0.785) and statistically significant ($t-statistic = 6.676$) slope estimate. The explanatory power associated with this regression $Adj\ R^2 = 61.6\%$ confirms that the monthly implied variance predicts the weekly realized variance.

Turning to the weekly implied variance, we can see that it predicts the future realized variance with a slope estimate of 0.713 ($t - statistic = 7.240$). The corresponding predictive power is equal to 64.9%. Several points are worth highlighting. First, the explanatory power of this regression model is higher than that of the model that relies on the monthly implied variance. This result indicates that the weekly implied variance has superior predictive ability for short-term risk than the often used monthly implied variance. Second, the slope estimates obtained from these univariate regressions are similar and significantly different from 1. While the earlier literature interpreted a slope that is significantly different from 1 as evidence of violation of the expectations hypothesis, Chernov (2007) and Prokopczuk and Wese Simen (2014) point out that this result can arise in a setting where the variance risk premium is time varying. It is thus interesting to analyze the average variance risk premium, defined as the difference between the implied variance and the contemporaneously estimated realized variance (Bollerslev et al., 2009), of each maturity. The penultimate row of Table 1 reports an average (annualized) variance risk premium of 2.280 % with a volatility of 4.659%. This average figure is very similar to that of the monthly variance swap (2.300%). Thus, there is evidence of a non-zero and time-varying variance risk premium in both the weekly and monthly implied variance indices.

3.1.b Multivariate Evidence

The preceding analysis shows that, when used alone, the weekly and monthly implied variance predict the future weekly realized variance. However, it does not directly shed light on the incremental information content of these predictors. To address this question, we include both maturities in the following encompassing regression model:

$$RV^w_{t+7} = \alpha + \beta IV^w_t + \gamma IV^m_t + \epsilon_{t+7} \tag{8}$$

where α is the intercept. β and γ are the slope parameters. All other variables are as previously defined.

The last row of Table 2 reports that the slope estimate associated with the weekly component remains significant ($t - stat = 2.195$), whereas that of the monthly implied variance loses its significance ($t - stat = 0.631$). Moreover, the explanatory power of the encompassing model ($Adj\ R^2 = 65.0\%$) is very similar to that of the univariate model which uses the weekly implied variance ($Adj\ R^2 = 64.9\%$) as a forecasting variable. Taken as a whole, the results suggest that the weekly implied variance drives out the information content of the monthly implied variance. This set of findings is important because empirical studies routinely discard short-term options data on the grounds that they are noisy and thus uninformative. Our results caution, that by following this approach, one throws away valuable information about short-term risk. These findings are also interesting given the growing practice of using the monthly implied variance to predict short-term realized variance. Our evidence reveals that this methodology may not be the best way of modeling the short-term realized variance.

3.2 Dissecting the Predictability

Having established the in-sample predictive power of the weekly implied variance series for realized variance, we now seek to further explore the channel through which this predictability arises.

3.2.a Framework

Our starting point is the theory of quadratic variation (Barndorff-Nielsen and Shephard, 2002) which posits that the realized variance of an asset return can be decomposed into components linked to the (i) continuous variation and the (ii) jump variation of the asset returns. More formally, we have:

$$RV^{w}_{t+7} = CV^{w}_{t+7} + JV^{w}_{t+7} \tag{9}$$

where CV^{w}_{t+7} and JV^{w}_{t+7} are the weekly continuous and jump variations of the asset returns computed over the week ending on day $t+7$, respectively.

This insight suggests that there are two channels through which short-term implied variance may be informative about next week's realized variance. The first possibility is that the weekly implied variance contains information about the continuous variation of returns. The second possibility is that the weekly implied variance is informative about the jump variation.

Barndorff-Nielsen and Shephard (2002) propose the bipower variation as an estimator of the continuous variation of asset returns. Andersen et al. (2012) subsequently establish that the $MedRV$ estimator has better properties than the bipower variation. Thus, we use the $MedRV$ estimator of the continuous variation of returns:

$$CV^{w}_{t+7} = \underbrace{\frac{52 \times m\pi}{(6-4\sqrt{3}+\pi)(m-2)} \sum_{j=0}^{N^{w}_{t+7}-1} \sum_{k=3}^{m} \text{median}(|r_{t+7-j,k}|, |r_{t+7-j,k-1}|, |r_{t+7-j,k-2}|)^2}_{MedRV\ \ Estimator} \tag{10}$$

where median$(\cdot)$ is the median operator. All other variables are as previously defined.

Re-arranging Equation (9), it is straightforward to extract the jump variation component:[13]

$$JV^w_{t+7} = RV^w_{t+7} - CV^w_{t+7} \tag{11}$$

3.2.b Continuous Variation

We regress the time series of the continuous variation on a constant and the lagged implied variance series:

$$CV^w_{t+7} = \alpha + \beta IV^x_t + \epsilon_{t+7} \tag{12}$$

where all variables are as previously defined.

Panel A of Table 3 shows that, in univariate regressions, each maturity of the implied variance predicts the continuous component of the realized variance. This conclusion is borne out by the significant slope estimates in univariate regressions. Similar to our analysis of the realized variance, we note that the weekly implied variance boasts the higher predictive power ($Adj\ R^2 = 62.0\%$) of the two variables. Combining the two implied variance series in the encompassing model yields an $Adj\ R^2$ of 62.1% that is very close to that of the univariate model based on the weekly implied variance alone. Moreover, the weekly implied variance drives out the significance of the monthly implied variance. This result echoes our earlier conclusion.

[13]As an additional check, we implement the jump tests used in Andersen et al. (2007) to identify significant jumps. If there are no significant jumps on a given day, then the continuous variation is equal to the realized variance and the jump variation takes the value 0. Otherwise, the continuous variation corresponds to the estimate given by the $MedRV$ estimator and the jump variation is the difference between the realized and continuous variations. The decomposition results are qualitatively similar to those of our benchmark approach. As a result, we do not tabulate these findings.

3.2.c Jump Variation

We now estimate the following forecasting regression:

$$JV_{t+7}^{w} = \alpha + \beta IV_{t}^{x} + \epsilon_{t+7} \tag{13}$$

where all variables are as previously defined.

Panel B of Table 3 documents that it is harder to accurately model the jump variation than the continuous variation. This conclusion is evidenced by the lower explanatory power for the jump variation compared to that of the continuous variation (see Panel A).[14] We can see that each maturity of the implied variance individually predicts the weekly jump variation. Additionally, the weekly implied variance subsumes the information content of the monthly implied variance.

4 Implied vs. Time Series Models

The previous section shows that the weekly implied variance is superior to the monthly implied variance in-sample. However, it is not clear how it compares to other sophisticated time series models that have been recently discussed in the literature. This section starts by presenting the competing models. Next, it assesses their in- and out-of-sample performance.

[14]Using various time series models, Busch et al. (2011) document a similar result at the monthly horizon.

4.1 Introducing the Competing Models

HAR We use the heterogeneous autoregressive (HAR) realized variance model (Corsi, 2009) as our benchmark:

$$RV^w_{t+7} = \alpha + \beta RV^d_t + \gamma RV^w_t + \delta RV^m_t + \epsilon_{t+7} \tag{14}$$

where α is the intercept. β, γ and δ are the slope parameters. RV^d_t and RV^m_t are the (annualized) daily and monthly realized variance at time t, respectively. These series are computed using the estimators below:

$$RV^d_t = 252 \times \sum_{k=1}^{m} r^2_{t,k} \tag{15}$$

$$RV^m_t = 12 \times \sum_{j=0}^{N^m_t - 1} \sum_{k=1}^{m} r^2_{t-j,k} \tag{16}$$

where the numbers 252 and 12 serve to annualize the daily and monthly realized variance estimates. N^m_t is the number of trading days in the calendar month ending on day t.

CHAR Building on the work of Andersen et al. (2007), we also consider the continuous heterogeneous autoregressive (CHAR) realized variance model that replaces each historical variance in the HAR model with the continuous variation of the corresponding maturity:

$$RV^w_{t+7} = \alpha + \beta CV^d_t + \gamma CV^w_t + \delta CV^m_t + \epsilon_{t+7} \tag{17}$$

where CV_t^d and CV_t^m are the (annualized) daily and monthly continuous variations at time t, respectively:

$$CV_t^d = \frac{252 \times m\pi}{(6 - 4\sqrt{3} + \pi)(m-2)} \sum_{k=3}^{m} \text{median}(|r_{t,k}|, |r_{t,k-1}|, |r_{t,k-2}|)^2 \quad (18)$$

$$CV_t^m = \frac{12 \times m\pi}{(6 - 4\sqrt{3} + \pi)(m-2)} \sum_{j=0}^{N_t^m - 1} \sum_{k=3}^{m} \text{median}(|r_{t-j,k}|, |r_{t-j,k-1}|, |r_{t-j,k-2}|)^2 \quad (19)$$

HAR–J We also analyze the performance of the HAR–J model, which augments the HAR model with the lagged significant daily jump variation as forecasting variable:

$$RV_{t+7}^w = \alpha + \beta RV_t^d + \gamma RV_t^w + \delta RV_t^m + \eta J_t^d + \epsilon_{t+7} \quad (20)$$

where α, β, γ, δ and η are parameters to estimate. J_t^d is the (annualized) statistically significant daily jump variation at time t. To obtain this quantity, we modify the test statistic presented in Huang and Tauchen (2005) that relies on the bipower variation and the realized quarticity to take advantage of the more robust estimators of the continuous variation ($MedRV$) and realized quarticity ($MedRQ$):

$$J_t^d = I_{z_t^d > \phi_{1-\alpha}}(RV_t^d - CV_t^d) \quad (21)$$

$$z_t^d = m^{1/2} \left(\frac{\frac{RV_t^d - CV_t^d}{RV_t^d}}{\sqrt{0.96 \max(1, \frac{MedRQ_t^d}{(CV_t^d)^2})}} \right) \quad (22)$$

$$where \quad (23)$$

$$MedRQ_t^d = \frac{252^2 \times 3m^2\pi}{(72 - 52\sqrt{3} + 9\pi)(m-2)} \sum_{k=3}^{m} \text{median}(|r_{t,k}|, |r_{t,k-1}|, |r_{t,k-2}|)^4 \quad (24)$$

$\phi_{1-\alpha}$ is the critical value from the cumulative standard normal distribution at confidence level $1-\alpha$. I is the indicator function. Similar to Andersen et al. (2007), we employ $\alpha = 99.9\,\%$ throughout this paper.

HAR–C–J We also evaluate the forecasting performance of the HAR–C–J model (Andersen et al., 2007). Essentially, this model decomposes each historical variance in the HAR model into the corresponding significant jump variation and the corresponding continuous variation:

$$RV^w_{t+7} = \alpha + \beta C^d_t + \gamma J^d_t + \delta C^w_t + \eta J^w_t + \theta C^m_t + \kappa J^m_t + \epsilon_{t+7} \tag{25}$$

where α, β, γ, δ, η, θ and κ are parameters to estimate. J^w_t and J^m_t denote the (annualized) weekly and monthly significant jump variation at time t, respectively. C^d_t, C^w_t and C^m_t are the corresponding daily, weekly and monthly continuous variation at time t, respectively:

$$J^w_t = I_{z^w_t > \phi_{1-\alpha}}(RV^w_t - CV^w_t) \tag{26}$$

$$z^w_t = m^{1/2}\left(\frac{\frac{RV^w_t - CV^w_t}{RV^w_t}}{\sqrt{0.96\max(1, \frac{MedRQ^w_t}{(CV^w_t)^2})}}\right) \tag{27}$$

$$MedRQ^w_t = A_1 \times \sum_{j=0}^{N^w_t - 1} \sum_{k=3}^{m} \text{median}(|r_{t-j,k}|, |r_{t-j,k-1}|, |r_{t-j,k-2}|)^4 \tag{28}$$

$$A_1 = \frac{52^2 \times 3m^2\pi}{(72 - 52\sqrt{3} + 9\pi)(m-2)} \tag{29}$$

$$C^w_t = I_{z^w_t > \phi_{1-\alpha}} CV^w_t + I_{z^w_t \leq \phi_{1-\alpha}} RV^w_t \tag{30}$$

$$J^m_t = I_{z^m_t > \phi_{1-\alpha}}(RV^m_t - CV^m_t) \tag{31}$$

$$z^m_t = m^{1/2}\left(\frac{\frac{RV^m_t - CV^m_t}{RV^m_t}}{\sqrt{0.96\max(1, \frac{MedRQ^m_t}{(CV^m_t)^2})}}\right) \tag{32}$$

$$MedRQ_t^m = A_2 \times \sum_{j=0}^{N_t^m-1} \sum_{k=3}^{m} \text{median}(|r_{t-j,k}|, |r_{t-j,k-1}|, |r_{t-j,k-2}|)^4 \quad (33)$$

$$A_2 = \frac{12^2 \times 3m^2\pi}{(72 - 52\sqrt{3} + 9\pi)(m-2)} \quad (34)$$

$$C_t^m = I_{z_t^m > \phi_{1-\alpha}} CV_t^m + I_{z_t^m \leq \phi_{1-\alpha}} RV_t^m \quad (35)$$

$$(36)$$

SHAR Patton and Sheppard (2015) document the good empirical performance of the semi-variance heterogeneous autoregressive model (SHAR) model. Essentially, this model modifies the HAR specification by decomposing each historical variance term into positive and negative semivariance components:[15]

$$RV_{t+7}^w = \alpha + \beta SV_t^{d+} + \gamma SV_t^{d-} + \delta SV_t^{w+} + \eta SV_t^{w-} + \theta SV_t^{m+} + \kappa SV_t^{m-} + \epsilon_{t+7} \quad (37)$$

where α, β, γ, δ, η, θ and κ are parameters to estimate. SV_t^{d+}, SV_t^{w+} and SV_t^{m+} are the positive (annualized) daily, weekly and monthly semivariances at time t, respectively. SV_t^{d-}, SV_t^{w-} and SV_t^{m-} are the negative (annualized) daily, weekly and monthly semivariances at time t, respectively. We compute these variables below:

$$SV_t^{d+} = 252 \times \sum_{k=1}^{m} r_{t,k}^2 I_{r_{t,k}>0} \quad (38)$$

$$SV_t^{d-} = 252 \times \sum_{k=1}^{m} r_{t,k}^2 I_{r_{t,k}<0} \quad (39)$$

$$SV_t^{w+} = 52 \times \sum_{j=0}^{N_t^w-1} \sum_{k=1}^{m} r_{t-j,k}^2 I_{r_{t-j,k}>0} \quad (40)$$

$$SV_t^{w-} = 52 \times \sum_{j=0}^{N_t^w-1} \sum_{k=1}^{m} r_{t-j,k}^2 I_{r_{t-j,k}<0} \quad (41)$$

[15]We have also considered more parsimonious specifications of the SHAR where we only decompose the historical variance of a specific horizon into positive and negative semivariance components. We found very little to distinguish between these alternative specifications.

$$SV_t^{m+} = 12 \times \sum_{j=0}^{N_t^m-1} \sum_{k=1}^{m} r_{t-j,k}^2 I_{r_{t-j,k}>0} \qquad (42)$$

$$SV_t^{m-} = 12 \times \sum_{j=0}^{N_t^m-1} \sum_{k=1}^{m} r_{t-j,k}^2 I_{r_{t-j,k}<0} \qquad (43)$$

4.2 In-Sample Evidence

Table 4 sheds light on the in-sample forecasting performance of each model. Starting with the HAR model, we can see that it yields an explanatory power of 59.8 %. The CHAR model yields a comparable explanatory power of 59.6 %. The fit of the HAR–J model ($Adj\ R^2 = 59.8\,\%$) to the data is similar to that of the HAR. This result arises because the exposure to the significant daily jump variation is not statistically significant. Turning to the HAR–C–J and SHAR models, we can see that they slightly improve on the benchmark HAR. This result is evidenced by the $Adj\ R^2$ of 60.8 % and 60.7 % for the HAR–C–J and SHAR models, respectively. This conclusion is consistent with the in-sample finding of Patton and Sheppard (2015).

Comparing the $Adj\ R^2$ of Tables 2 and 4, we can see that the forecasting model which uses the weekly implied variance as sole predictor achieves the highest explanatory power ($Adj\ R^2 = 64.9\,\%$). This result leads us to the conclusion that the weekly implied variance performs better than the HAR model and its extensions in-sample.

4.3 Out-of-Sample Evidence

We next investigate whether the in-sample predictability results also extend out-of-sample. We use a rolling window containing 4 years of daily data to estimate the forecasting models in Equations (7), (14), (17), (20), (25) and (37).[16] We then use all

[16] A window of 4 years is consistent with the research of Bollerslev et al. (2016) and Patton and Sheppard (2015).

relevant information available in real-time to generate the conditional expectation of next week's realized variance. Similar to Bollerslev et al. (2016), we subject these forecasts to the "insanity" filter to guard ourselves against implausible variance forecasts. If the forecast is higher (lower) than the highest (lowest) weekly realized variance observed in the estimation window, we set the forecast to the average weekly realized variance in the estimation window. This filter also enables us to avoid the situation where the variance forecast could be negative.[17]

Repeating the steps above for each rolling window, we obtain the time series of the out-of-sample variance forecasts which we then compare to the realized variance. We compute the mean squared error (MSE) and ($QLIKE$) loss functions below:

$$MSE = \frac{1}{T}\sum_{t=1}^{T}\left(RV_{t+7}^{w} - E_t(RV_{t+7}^{w})\right)^2 \tag{44}$$

$$QLIKE = \frac{1}{T}\sum_{t=1}^{T}\left(\frac{RV_{t+7}^{w}}{E_t(RV_{t+7}^{w})} - \log\frac{RV_{t+7}^{w}}{E_t(RV_{t+7}^{w})} - 1\right) \tag{45}$$

where T is the total number of out-of-sample forecasts. $E_t(RV_{t+7}^{w})$ is the expectation at time t of the variance to be realized at $t+7$. All other variables are as previously defined.

Patton (2011) shows that these loss functions are robust to the noise in the realized variance proxy, making them well-suited for our analysis.[18] Table 5 reports the ratio of the loss function [name in row] associated with the model [name in column] over that of the benchmark HAR model. An entry equal to 1 indicates that the model [name in column] does as well as the benchmark HAR model. Entries lower than 1 suggest that the model [name in column] achieves lower average forecasting error than the HAR model. Conversely, entries that are greater than 1 indicate that the forecast errors of the model

[17] As a robustness check, we remove the filter and obtain similar results. These findings are available upon request.

[18] Section *5.1* discusses further the implications of measurement errors.

[name in column] are higher than those of the HAR model.

Focusing on the entries reported under the header "IV^w", we can see that the MSE and $QLIKE$ ratios are equal to 0.797 and 0.870, respectively. This set of numbers reveals that a forecasting model based on the weekly implied variance reduces the forecasting errors of the benchmark HAR by 20.3 % (MSE) and 13.0 % ($QLIKE$). Looking at all the competing models, we can see that their performance is generally inferior to that of the weekly implied variance.

We formally test whether the null hypothesis that the average forecast error associated with the weekly implied variance is equal to or greater than that of the best forecasting model among its competitors (IV^m, HAR, $CHAR$, $HAR-J$, $HAR-C-J$ and $SHAR$). The alternate hypothesis is that the weekly implied variance delivers smaller forecast errors than its competitors. To implement this test, we modify the reality check test of White (2000) as in Bollerslev et al. (2016). In our empirical implementation, we use the stationary bootstrap of Politis and Romano (1994) with 9,999 re-samplings and an average block length of 10.[19] The null hypothesis is rejected at the 5% significance level with p-values of 1.9% and 1.6% for the MSE and $QLIKE$ loss functions, respectively. We conclude that the weekly implied variance achieves significantly lower forecasting errors compared to its competitors.

5 What About ...

This section presents several robustness checks. First, we explore whether our results are affected by measurement errors in the historical variance estimates. Second, we study the impact of potential errors in the method of interpolation used to obtain the implied

[19] We experimented with different lengths of the block and obtained very similar results.

variance series. Third, we assess the sensitivity of our findings to the estimator of implied variance.

5.1 Noise in Realized Variance?

Bollerslev et al. (2016) recently propose a forecasting model that extends the HAR model by taking into account the measurement errors in the historical variance. These errors arise from the fact that historical variance, the main input to the HAR model, is not directly observable. Thus, one needs to estimate the historical variance before using it for forecasting, leading to the errors-in-variables problem. The authors introduce a modeling framework, termed HAR–RQ, that aims to capture the heteroskedasticity of the measurement errors and improve the realized variance forecasts:[20]

$$
\begin{aligned}
RV^w_{t+7} &= \alpha + (\beta + \gamma\sqrt{MedRQ^d_t})RV^d_t + (\delta + \eta\sqrt{MedRQ^w_t})RV^w_t \\
&\quad + (\theta + \kappa\sqrt{MedRQ^m_t})RV^m_t + \epsilon_{t+7} \qquad (46)
\end{aligned}
$$

where α is the intercept. β, γ, δ, η, θ and κ are slope parameters.[21]

Using the same out-of-sample methodology as before, we analyze the performance of this model out-of-sample. Consistent with Bollerslev et al. (2016), Table 5 shows that this model achieves lower forecasting errors than the CHAR, HAR–J and HAR–C–J models. However, its performance is weaker than that of the weekly implied variance, indicating that short-term implied variance provides more accurate variance forecasts

[20] Bollerslev et al. (2016) use a square root specification for the measurement error correction on the basis that it has an imbued robustness. We also experimented with the log specification and found it to deliver inferior forecasting performance compared to the square root specification. This finding is consistent with the authors' argument and their own empirical results.

[21] Bollerslev et al. (2016) use the realized quarticity as defined in Barndorff-Nielsen (2002) rather than the more robust MedRQ estimator of Andersen et al. (2012). We have also repeated the analysis using the same estimator as the authors. The results were qualitatively similar.

than the HAR–RQ model. Our untabulated analysis also reveals that the difference in the performance of the two models is statistically significant.

In addition, we considered more parsimonious specifications of the HAR–RQ model in Equation (46): the HAR–RQ–D (Equation(47)), HAR–RQ–W (Equation(48)) and HAR–RQ–M (Equation(49)).

$$RV^w_{t+7} = \alpha + (\beta + \gamma\sqrt{MedRQ^d_t})RV^d_t + \delta RV^w_t + \theta RV^m_t + \epsilon_{t+7} \quad (47)$$

$$RV^w_{t+7} = \alpha + \beta RV^d_t + (\delta + \eta\sqrt{MedRQ^w_t})RV^w_t + \theta RV^m_t + \epsilon_{t+7} \quad (48)$$

$$RV^w_{t+7} = \alpha + \beta RV^d_t + \delta RV^w_t + (\theta + \kappa\sqrt{MedRQ^m_t})RV^m_t + \epsilon_{t+7} \quad (49)$$

Generally, our untabulated analysis reveals that these parameterizations do not beat the weekly implied variance. For instance, the HAR–RQ–D and the HAR–RQ–W models yield MSE (QLIKE) loss ratios of 0.864 (0.917) and 0.861 (0.902), respectively. While the HAR–RQ–W performs better than the more general HAR–RQ specification in Equation (46), a finding consistent with the work of Bollerslev et al. (2016), it does not outperform the weekly implied variance.

From a theoretical standpoint, the noise in the realized variance series should be larger at lower sampling frequencies (Barndorff-Nielsen, 2002). Thus, by sampling data at finer frequencies, one should be able to dampen the effect of the noise. This insight motivates us to increase the sampling frequency from 5-minute to 1-minute and repeat our out-of-sample analysis. Table 6 points to the same conclusion as Table 5: the weekly implied variance yields the biggest improvement in forecast accuracy.

5.2 The Method of Interpolation?

Our implied variance is based on a constant maturity series. To do this, one needs to interpolate the implied variance across maturities, making the analysis dependent on the interpolation method. Thus, it is interesting to assess the sensitivity of our results to the method of interpolation.

We consider an alternative approach where in order to obtain the implied variance of the weekly (monthly) maturity, we use the variance implied by option contracts of maturity closest to 1 week (month). Essentially, this approach is the nearest neighborhood interpolation method. With this new time series, we repeat our out-of-sample analysis. Overall, Table 7 confirms our main results. The weekly implied variance is a strong predictor of short-term risk as it outperforms the monthly implied variance as well as the recently proposed variance forecasting models.

5.3 The Implied Variance Estimator?

Our interest in the Bakshi et al. (2003) estimator is motivated by its robustness to jumps (Du and Kapadia, 2013). However, most existing studies, e.g. Jiang and Tian (2005) and Taylor et al. (2010), use the Britten-Jones and Neuberger (2000) estimator, making our results difficult to directly compare to those of the extant literature. We use the numerical scheme presented in Section *2.1.b* to implement the Britten-Jones and

Neuberger (2000) estimator of implied variance:[22]

$$IV_t^{\tau} = \frac{2e^{\mathrm{rf}_t \frac{\tau}{360}}}{\frac{\tau}{360}} \left(\int_0^{S_t} \frac{P_t(\tau, K)}{K^2} dK + \int_{S_t}^{\infty} \frac{C_t(\tau, K)}{K^2} dK \right) \tag{50}$$

where all variables are as previously discussed.

We use the resulting time series to repeat our out-of-sample analysis that we summarize in Table 8. The loss ratios associated with the implied variance series are generally higher than those based on the Bakshi et al. (2003) estimator (see Table 5). This is not surprising since the Britten-Jones and Neuberger (2000) is biased in the presence of jumps (Du and Kapadia, 2013), resulting in larger forecast errors. Additionally, the weekly implied variance series outperforms its monthly counterpart. This is true for both the MSE and the $QLIKE$ loss functions. Furthermore, the weekly implied variance series outperforms all its competitors. Overall, these results are consistent with our main findings.

6 Conclusion

We exploit the launch of weekly options on the S&P 500 index to analyze the information content of short-term options for short-term realized variance. Our results reveal that the weekly implied variance is a powerful predictor of future weekly realized variance. Dissecting the realized variance into continuous and jump components, we find that the

[22]Generally, the literature directly uses the squared value of the volatility index (VIX), instead of computing the monthly implied variance series using the Britten-Jones and Neuberger (2000) formula and the numerical method presented in Section *2.1.b*. It is interesting to compare the performance of these two approaches. By doing so, one can shed light on the impact of the numerical method on the results. Untabulated results show that the time series of the square of the VIX index and our own Britten-Jones and Neuberger (2000) series of monthly maturity are highly correlated and very similar. Empirically, the squared VIX series yields a $QLIKE$ loss ratio of 1.251 that is comparable to the 1.217 figure based on our computation of the monthly implied variance using the Britten-Jones and Neuberger (2000) estimator (see Table 8). We thus conclude that the numerical method does not have a major impact on the variance forecasting results.

weekly implied variance strongly predicts both components.

The information content of the short-term implied variance subsumes that of the monthly implied variance, suggesting that it is informationally more efficient than the often used monthly implied variance. Out-of-sample, the weekly implied variance outperforms not only the monthly implied variance but also the HAR model and its recent extensions.

Our evidence carries implications for both academics and practitioners. For practitioners, it would be interesting for the Chicago Board of Options Exchange (CBOE) to compute and disseminate the time series of the weekly volatility index. This series, which would sit alongside the popular 30-day volatility index, would be useful for market participants to better gauge and manage short-term risk. For academics, our results suggest that one would benefit from not discarding short-term options on the grounds that they are illiquid and thus uninformative. Our analysis clearly shows that these options are informative and crowd out the information content of monthly options that have been analyzed hitherto.

References

Andersen, T. G. and Bollerslev, T. (1998). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review*, 39(4):885–905.

Andersen, T. G., Bollerslev, T., and Diebold, F. X. (2007). Roughing it up: Including jump components in the measurement, modeling, and forecasting of return volatility. *Review of Economics and Statistics*, 89(4):701–720.

Andersen, T. G., Bollerslev, T., Diebold, F. X., and Labys, P. (2003). Modelling and forecasting realized volatility. *Econometrica*, 71(2):579–625.

Andersen, T. G., Dobrev, D., and Schaumburg, E. (2012). Jump-robust volatility estimation using nearest neighbor truncation. *Journal of Econometrics*, 169(1):75–93.

Andersen, T. G., Fusari, N., and Todorov, V. (2017). Short-term market risk implied by weekly options. *Journal of Finance*, 72(3):1335–1386.

Bakshi, G., Kapadia, N., and Madan, D. (2003). Stock return characteristics, skew laws, and the differential pricing of individual equity options. *Review of Financial Studies*, 16(1):101–143.

Barndorff-Nielsen, O. E. (2002). 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. and Shephard, N. (2002). Estimating quadratic variation using realized variance. *Journal of Applied Econometrics*, 17(5):457–477.

Bekaert, G. and Hoerova, M. (2014). The VIX, the variance premium and stock market volatility. *Journal of Econometrics*, 183(2):181–192.

Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654.

Blair, B. J., Poon, S.-H., and 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.

Bollerslev, T., Patton, A. J., and Quaedvlieg, R. (2016). Exploiting the errors: A simple approach for improved volatility forecasting. *Journal of Econometrics*, 192(1):1–18.

Bollerslev, T., Tauchen, G., and Zhou, H. (2009). Expected stock returns and variance risk premia. *Review of Financial Studies*, 22(11):4463–4492.

Bollerslev, T. and Todorov, V. (2011). Tails, fears, and risk premia. *Journal of Finance*, 66(6):2165–2211.

Britten-Jones, M. and Neuberger, A. (2000). Option prices, implied price processes and stochastic volatility. *Journal of Finance*, 55(2):839–866.

Busch, T., Christensen, B. J., and 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.

Carr, P. and Wu, L. (2009). Variance risk premiums. *Review of Financial Studies*, 22(3):1311–1341.

Chernov, M. (2007). On the role of risk premia in volatility forecasting. *Journal of Business & Economic Statistics*, 25(4):411–426.

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

Drechsler, I. and Yaron, A. (2011). What's vol got to do with it. *Review of Financial Studies*, 24(1):1–45.

Du, J. and Kapadia, N. (2013). The tail in the volatility index. *EFA Meetings 2013 Paper*.

Huang, X. and Tauchen, G. (2005). The relative contribution of jumps to total price variance. *Journal of Financial Econometrics*, 3(4):456–499.

Jiang, G. J. and Tian, Y. S. (2005). The model-free implied volatility and its information content. *Review of Financial Studies*, 18(4):1305–1342.

Kourtis, A., Markellos, R. N., and Symeonidis, L. (2016). An international comparison of implied, realized, and GARCH volatility forecasts. *Journal of Futures Markets*, 36(12):1164–1193.

Lee, S. S. and Mykland, P. A. (2008). Jumps in financial markets: A new nonparametric test and jump dynamics. *Review of Financial Studies*, 21(6):2535–2563.

Mincer, J. A. and Zarnowitz, V. (1969). The evaluation of economic forecasts. In Mincer, J., editor, *Economic Forecasts and Expectations: Analysis of Forecasting Behavior and Performance*, pages 1–46. Cambridge, MA: Elsevier.

Newey, W. K. and West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708.

Patton, A. J. (2011). Volatility forecast comparison using imperfect volatility proxies. *Journal of Econometrics*, 160(1):246–256.

Patton, A. J. and Sheppard, K. (2015). Good volatility, bad volatility: Signed jumps and the persistence of volatility. *Review of Economics and Statistics*, 97(3):683–697.

Politis, D. N. and Romano, J. P. (1994). The stationary bootstrap. *Journal of the American Statistical Association*, 89(428):1303–1313.

Pong, S., Shackleton, M. B., Taylor, S. J., and 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. and Granger, C. W. (2003). Forecasting volatility in financial markets: A review. *Journal of Economic Literature*, 41(2):478–539.

Prokopczuk, M. and Wese Simen, C. (2014). The importance of the volatility risk premium for volatility forecasting. *Journal of Banking & Finance*, 40:303–320.

Prokopczuk, M. and Wese Simen, C. (2016). What makes the market jump? *EFA Meetings 2016 Paper.*

Taylor, S. J., Yadav, P. K., and Zhang, Y. (2010). The information content of implied volatilities and model-free volatility expectations: Evidence from options written on individual stocks. *Journal of Banking & Finance*, 34(4):871–881.

White, H. (2000). A reality check for data snooping. *Econometrica*, 68(5):1097–1126.

Table 1: Descriptive Statistics

This table presents key summary statistics. IV^w and IV^m denote the (annualized) implied variance of weekly and monthly horizons, respectively. RV^d, RV^w and RV^m denote the (annualized) realized variance of daily, weekly and monthly horizons, respectively. VRP^w and VRP^m are the weekly and monthly variance risk premia, respectively. Similar to Bollerslev et al. (2009), the variance risk premium of a given maturity is defined as the difference between the implied variance of that maturity and the contemporaneously computed realized variance of the same maturity. Returns data are sampled at the 5-minute frequency. $Mean$ is the average value of the daily time series of the variable [name in row]. Std, $Skew$ and $Kurt$ denote the standard deviation, skewness and kurtosis of the variable [name in row], respectively.

	$Mean$	Std	$Skew$	$Kurt$
IV^w	5.960%	8.827%	4.385	27.976
IV^m	5.977%	7.815%	3.884	21.890
RV^d	3.609%	7.966%	6.471	57.387
RV^w	3.680%	7.574%	6.004	51.738
RV^m	3.677%	6.842%	4.756	29.075
VRP^w	2.280%	4.659%	0.506	40.845
VRP^m	2.300%	3.107%	1.109	29.283

Table 2: In-Sample Results: Implied Variance

This table summarizes the results of regressions of the daily time series of the (annualized) weekly realized variance on a constant and the lagged forecasting variable(s) [name in column]. α denotes the intercept parameter. IV^w and IV^m are the Bakshi et al. (2003) weekly and monthly implied variances, respectively. We present in parentheses the Newey and West (1987) corrected t-statistics with 10 lags. $Adj\ R^2$ is the adjusted R-squared of the regression model. Returns are sampled at the 5-minute frequency.

α	IV^w	IV^m	$Adj\ R^2$
-0.005	0.713		0.649
(-1.376)	(7.240)		
-0.010		0.785	0.616
(-2.140)		(6.676)	
-0.007	0.597	0.137	0.650
(-2.393)	(2.195)	(0.631)	

Table 3: Continuous vs. Jump Variation

This table dissects the source of the predictability of the weekly realized variance. Panel A presents the results of regressions of the daily time series of the (annualized) weekly continuous variation, estimated using the $MedRV$ estimator of Andersen et al. (2012), on a constant and the lagged variable(s) [name in column]. Panel B summarizes the results of the regression of the daily time series of the (annualized) weekly jump variation on a constant and the forecasting variable(s) [name in column]. α denotes the intercept parameter. IV^w and IV^m are the Bakshi et al. (2003) weekly and monthly implied variances, respectively. We present in parentheses the Newey–West corrected t-statistics with 10 lags. $Adj\ R^2$ is the adjusted R-squared of the regression. Returns are sampled at the 5-minute frequency.

Panel A: Continuous Variation

α	IV^w	IV^m	$Adj\ R^2$
-0.007	0.635		0.620
(-1.922)	(6.632)		
-0.011		0.697	0.586
(-2.625)		(6.321)	
-0.008	0.545	0.106	0.621
(-3.100)	(2.070)	(0.501)	

Panel B: Jump Variation

α	IV^w	IV^m	$Adj\ R^2$
0.002	0.078		0.473
(3.503)	(7.310)		
0.001		0.087	0.461
(1.765)		(5.769)	
0.002	0.052	0.031	0.477
(2.438)	(2.808)	(1.467)	

Table 4: In-Sample Results: Other Models

This table summarizes the results of regressions of the daily time series of the (annualized) weekly realized variance on a constant and the lagged forecasting variable(s) [name in row]. α denotes the intercept parameter. RV^d, RV^w and RV^m denote the (annualized) daily, weekly and monthly realized variance series, respectively. C and J indicate the (annualized) continuous and (annualized) discontinuous variations, respectively. The associated superscripts indicate that we compute these quantities for the daily (d), weekly (w) and monthly (m) horizons, respectively. SV^{d+} and SV^{d-} are the (annualized) positive and negative daily semivariances, respectively. SV^{w+} and SV^{w-} denote the (annualized) positive and negative weekly semivariances, respectively. SV^{m+} and SV^{m-} are the (annualized) positive and negative monthly semivariances, respectively. We present in parentheses the Newey and West (1987) corrected t-statistics with 10 lags. $Adj\ R^2$ is the adjusted R-squared of the regression. Returns are sampled at the 5-minute frequency.

	HAR	$CHAR$	$HAR-J$	$HAR-C-J$	$SHAR$
α	0.005	0.007	0.005	0.002	0.006
	(2.592)	(3.271)	(2.816)	(0.464)	(2.620)
RV^d	0.342		0.343		
	(2.650)		(2.607)		
RV^w	0.216		0.215		
	(3.254)		(3.214)		
RV^m	0.311		0.310		
	(3.679)		(3.669)		
C^d		0.340		0.338	
		(2.603)		(2.705)	
C^w		0.229		0.219	
		(3.387)		(2.938)	
C^m		0.299		0.298	
		(3.514)		(3.483)	
J^d			-0.084	0.443	
			(-0.397)	(3.070)	
J^w				-1.911	
				(-1.255)	
J^M				5.292	
				(1.297)	
SV^{d+}					0.075
					(0.704)
SV^{d-}					0.482
					(2.851)
SV^{w+}					-0.081
					(-0.142)
SV^{w-}					0.577
					(1.085)
SV^{m+}					0.084
					(0.076)
SV^{m-}					0.546
					(0.543)
$Adj\ R^2$	0.598	0.596	0.598	0.608	0.607

Table 5: Out-of-Sample Results

This table presents the ratio of the average loss based on the function [name in row] of the model [name in column] over that of the HAR model. We use a rolling window of 4 years to estimate the parameters of the forecasting models. HAR denotes the forecasting model that uses the daily, weekly and monthly lagged realized variance series to predict the future weekly realized variance. IV^w and IV^m are the Bakshi et al. (2003) weekly and monthly implied variances, respectively. $CHAR$ is the continuous heterogeneous autoregressive model. $HAR-J$ extends the HAR model by including the lagged daily jump variation. The $HAR-C-J$ model decomposes the historical variances in the HAR into their continuous and jump components. The $SHAR$ model extends the HAR by splitting the lagged historical variance terms into the corresponding positive and negative semivariance components. $HAR-RQ$ takes into account the heteroskedasticity of the measurement error in all three maturities of lagged realized variance. Similar to Bollerslev et al. (2016), we proxy the heteroskedasticity of the measurement error in the realized variance with the square root of the realized quarticity of corresponding maturity. Returns are sampled at the 5-minute frequency.

	IV^w	IV^m	$CHAR$	$HAR-J$	$HAR-C-J$	$SHAR$	$HAR-RQ$
MSE	0.797	0.955	1.012	1.001	0.998	0.887	0.836
$QLIKE$	0.870	1.049	1.022	1.000	1.016	0.951	0.939

Table 6: Out-of-Sample Results (1-Minute Sampling Frequency)

This table presents the ratio of the average loss based on the function [name in row] of the model [name in column] over that of the HAR model. We use a rolling window of 4 years to estimate the parameters of the forecasting models. HAR denotes the forecasting model that uses the daily, weekly and monthly lagged realized variance series to predict the future weekly realized variance. IV^w and IV^m are the Bakshi et al. (2003) weekly and monthly implied variances, respectively. $CHAR$ is the continuous heterogeneous autoregressive model. $HAR-J$ extends the HAR model by including the lagged daily jump variation. The $HAR-C-J$ model decomposes the historical variances in the HAR into their continuous and jump components. The $SHAR$ model extends the HAR by splitting the lagged historical variance terms into the corresponding positive and negative semivariance components. $HAR-RQ$ takes into account the heteroskedasticity of the measurement error in all three maturities of lagged realized variance. Similar to Bollerslev et al. (2016), we proxy the heteroskedasticity of the measurement error in the realized variance with the square root of the realized quarticity of corresponding maturity. Returns are sampled at the 1-minute frequency.

	IV^w	IV^m	$CHAR$	$HAR-J$	$HAR-C-J$	$SHAR$	$HAR-RQ$
MSE	0.791	1.013	0.936	1.014	1.002	1.021	0.883
$QLIKE$	0.897	1.145	1.159	1.024	0.970	0.948	0.962

Table 7: Out-of-Sample Results (Nearest)

This table presents the ratio of the average loss based on the function [name in row] of the model [name in column] over that of the HAR model. We use a rolling window of 4 years to estimate the parameters of the forecasting models. HAR denotes the forecasting model that uses the daily, weekly and monthly lagged realized variance series to predict the future weekly realized variance. IV^w and IV^m are the Bakshi et al. (2003) weekly and monthly implied variances computed using options of maturities nearest to the weekly and monthly horizons, respectively. $CHAR$ is the continuous heterogeneous autoregressive model. $HAR-J$ extends the HAR model by including the lagged daily jump variation. in the HAR into their continuous and jump components. The $SHAR$ model extends the HAR by splitting the lagged historical variance terms into the corresponding positive and negative semivariance components. $HAR-RQ$ takes into account the heteroskedasticity of the measurement error in all three maturities of lagged realized variance. Similar to Bollerslev et al. (2016), we proxy the heteroskedasticity of the measurement error in the realized variance with the square root of the realized quarticity of corresponding maturity. Returns are sampled at the 5-minute frequency.

	IV^w	IV^m	$CHAR$	$HAR-J$	$HAR-C-J$	$SHAR$	$HAR-RQ$
MSE	0.763	0.861	1.012	1.001	0.998	0.887	0.836
$QLIKE$	0.811	0.915	1.022	1.000	1.016	0.951	0.939

Table 8: Out-of-Sample Results (Britten-Jones and Neuberger, 2000)

This table presents the ratio of the loss function [name in row] of the model [name in column] over that of the HAR model. We use a rolling window of 4 years to estimate the parameters of the forecasting models. HAR denotes the forecasting model that uses the daily, weekly and monthly lagged realized variance series to predict the future weekly realized variance. We compute the realized variance estimates with 1-minute return data. IV^w and IV^m are the Britten-Jones and Neuberger (2000) weekly and monthly implied variances, respectively. $CHAR$ is the continuous heterogeneous autoregressive model. $HAR-J$ extends the HAR model by including the lagged daily jump variation. The $HAR-C-J$ model decomposes the historical variances in the HAR into their continuous and jump components. The $SHAR$ model extends the HAR by splitting the lagged historical variance terms into the corresponding positive and negative semivariance components. $HAR-RQ$ takes into account the heteroskedasticity of the measurement error in all three maturities of lagged realized variance. Similar to Bollerslev et al. (2016), we proxy the heteroskedasticity of the measurement error in the realized variance with the square root of the realized quarticity of corresponding maturity. Returns are sampled at the 5-minute frequency.

	IV^w	IV^m	$CHAR$	$HAR-J$	$HAR-C-J$	$SHAR$	$HAR-RQ$
MSE	0.805	1.121	1.020	1.010	0.999	0.887	0.836
$QLIKE$	0.891	1.217	1.004	0.993	1.015	0.951	0.939