The Information Content of Short-Term Options

September 7, 2017

Abstract

Motivated by the growing trading activity in short-term (weekly) options, we propose the HAR–IV model that jointly uses the daily, weekly and monthly implied variance to predict realized variance. The HAR–IV model outperforms the HAR–RV model both in- and out-of-sample. An investor would pay up to 3.887 % per year to switch from the timing strategy based on the HAR–RV model to the strategy based on the HAR–IV model. Our results are robust to heteroscedastic measurement errors and several additional checks.

JEL classification: G11, G12

Keywords: Implied Variance, Predictability, Realized Variance, Term-Structure, Weekly Options

1 Introduction

In a recent study, Andersen et al. (2017) draw attention to an important development in the S&P 500 index option market. The authors show that weekly option contracts, which are short-term options, account for nearly 50 % of the total trading volume in the S&P 500 index options in 2015. This finding raises several questions about the information content of short-term options: Are short-term implied variance series informative about next-month's realized variance? If so, is this information content incremental to that of the monthly implied variance? What are the implications for forecasting models of realized variance? What is the economic value of this predictability? This study aims to answer these questions.

We make three contributions to the literature. First, we evaluate the information content of short-term implied variance for the monthly realized variance computed using 5-minute S&P 500 data. In a regression of monthly realized variance on a constant and the lagged daily implied variance, we obtain a statistically significant slope estimate and a high predictive power ($Adj\ R^2 = 0.743$). We thus conclude that the daily implied variance is informative about the future realized variance. In order to understand the channel through which this predictability arises, we decompose the realized variance into its continuous and jump components. We find that both channels are at work. The daily implied variance strongly predicts the continuous ($Adj\ R^2 = 0.718$) and jump ($Adj\ R^2 = 0.742$) variation. The results for the weekly implied variance series are similar to those of the daily implied variance series.

Second, we examine the extent to which the information content of the daily and weekly series is subsumed by the monthly implied variance. To this end, we propose the HAR–IV model that includes all three maturities, i.e. daily, weekly and monthly, to predict realized variance. The regression results suggest that the information content of the short-term implied variance series is not subsumed by the monthly implied variance series. The predictive power of the HAR–IV model ($Adj\ R^2 = 0.859$) is substantially higher than that afforded by the model based only on the monthly implied variance ($Adj\ R^2 = 0.635$).

Third, we compare the HAR–IV to the HAR–RV model of Corsi (2009). This analysis is interesting because, conceptually, the two models are similar in that they use information from the term-structure of variance to predict realized variance. To be more precise, the HAR–IV model uses the daily, weekly and monthly components of the term-structure of implied variance whereas the HAR–RV model uses the daily, weekly and monthly maturities of the term-structure of historical variance. We find that the in-sample predictive power of the HAR–IV is superior to that of the HAR–RV ($Adj\ R^2 = 0.775$). Similar conclusions emerge from the out-of-sample analysis. For instance, the mean squared percentage error (MSPE) of the HAR–IV model is 74 % lower than that of the HAR–RV model. The improvement in forecast accuracy is not only statistically significant but also economically meaningful. An agent with a risk aversion parameter equal to 3 would pay up to 3.887 % per year to switch from the timing strategy based on the HAR–RV model to the strategy based on the HAR–IV model.

One may be concerned about measurement errors. After all, weekly options are relatively new contracts and time-variations in their trading activity could introduce heteroscedastic measurement errors in the short-term implied variance series. To shed light on this, we propose a model with a time-varying slope parameter in the spirit of Bollerslev et al. (2016). In particular, we model the relationship between the realized variance and the lagged daily implied variance as a function of the illiquidity of the options that underpin the implied variance series. Empirically, we find that accounting for heteroscedastic measurement errors does not change our main conclusions. Our results are robust to several additional tests. We obtain similar findings using a sampling frequency of 1-minute to compute the realized (and historical) variance. Additionally, we show that our results are robust to concerns related to the asynchronous closing times between the spot and derivatives markets. We also consider alternative proxies for the illiquidity of derivative securities and reach similar conclusions. Finally, we consider a richer model that accounts for heteroscedastic measurement errors in all components of the HAR–IV model and reach the same findings.

Our research relates to the growing literature on short-term options. Bollerslev and Todorov (2011) and Andersen et al. (2015), among others, use options with maturity of less than 5 business days to learn about jumps. Andersen et al. (2017) are the first to provide an overview of weekly options. The authors exploit this new dataset to learn about the spot volatility and jump risk. Implicit in these studies is the assumption that short-term options are informative about the *state* of the underlying asset return process rather than the *expected* future variations in the volatility and jump intensity. We directly test this assumption by analyzing the predictability of next month's realized variance. We show that the implied variance of daily horizon is a strong predictor of the monthly realized variance.

Our study also adds to the broader literature on realized variance forecasting. Corsi (2009) proposes the HAR–RV model and documents its superior performance relative to the simple random walk model. Andersen et al. (2007) decompose the components of the HAR–RV into continuous and jump components. Patton and Sheppard (2015) propose an extension that separately uses positive and negative semi-variances. More recently, Bollerslev et al. (2016) extend the HAR–RV model to account for heteroscedastic measurement errors. We contribute to this strand of the literature by proposing the HAR–IV, which is the option-based analog of the HAR–RV. Empirically, we show that the HAR–IV model is superior to the HAR–RV model. We also propose a simple approach to extend the HAR–IV model to account for heteroscedastic measurement errors.

The remainder of this paper proceeds as follows. Section 2 introduces the methodology and the dataset. Section 3 discusses the performance of the HAR–IV model. Section 4 compares the HAR–IV and HAR–RV models. Section 5 discusses the effect of heteroscedastic measurement errors. Section 6 presents various robustness checks. Finally, Section 7 concludes.

2 Data and Methodology

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

2.1 Methodology

Realized Variance Our paper focuses on the predictability of 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) monthly realized variance as follows:

$$RV_{t+1}^{M} = 12 \times \sum_{j=0}^{N-1} \sum_{k=1}^{m} r_{t+1-\frac{j}{N},k}^{2} \tag{2}$$

where RV_{t+1}^{M} is the (annualized) monthly realized variance for the calendar month starting at time t and ending at time $t+1$. The number 12 indicates that the realized variance estimate is annualized. There are m intraday returns each trading day. N is the number of trading days in the month starting at time t and ending at $t+1$.[1]

Implied Variance We use the Bakshi et al. (2003) formula to compute the implied variance:[2]

$$IV_t = \frac{12}{\tau}\left(e^{\mathrm{rf_t}\frac{\tau}{12}}\mathrm{QUAD}_t - \mu_t^2\right) \tag{3}$$

where

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

$$\mu_t = e^{\mathrm{rf_t}\frac{\tau}{12}} - 1 - \frac{e^{\mathrm{rf_t}\frac{\tau}{12}}}{2}\mathrm{QUAD}_t \tag{5}$$

where IV_t is the (annualized) implied variance of time-to-maturity τ (expressed in months) observed at time t. $\mathrm{rf_t}$ is the τ-month (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.

[1] Obviously, the number of trading days could vary from one month to the next. For ease of notation, we simply suppress the time subscript on the variable N.

[2] Our interest in the Bakshi et al. (2003) implied variance rather than the Britten-Jones and Neuberger (2000) variance is motivated by the research of Du and Kapadia (2013), who show that the former is more robust to jumps. As a robustness check, we also consider the Britten-Jones and Neuberger (2000) variance. Our untabulated results point to qualitatively similar conclusions.

For ease of exposition, suppose we want to compute the implied variance of daily, weekly and monthly maturities on a given trading day. 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:

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

$$K_{t,U} = S_t e^{8\sigma_t} \quad (7)$$

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 (6) and (7). In practice, the strike prices available 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. We then use the trapezoidal rule to numerically evaluate the integrals in

Equation (4). Finally, we compute the implied variance as in Equation (3).[3]

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 daily (IV^D), weekly (IV^W) and monthly (IV^M) horizons. It is important to emphasize that we only interpolate between maturities. To be clear, we do not implement any extrapolation since this could introduce spurious spikes in the constant maturity implied variance series. For instance, if the shortest maturity of options available on that day is equal to 6 trading days, we compute the weekly and monthly implied variance but not the daily implied variance.

2.2 Data

We obtain high-frequency data on the S&P 500 index from Thomson Reuters Tick History (TRTH) to build the monthly 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 accurately measure realized variance. The dataset spans the period extending from January 1996 to August 2015. 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), we use a 5-minute sampling frequency.[4] At the end of each 5-minute interval, we use the most recent mid-quote to proxy for the closing price of that interval.

[3]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 is standard in the literature. Our choice of 8 standard deviations is consistent with earlier work, e.g. Carr and Wu (2009). Jiang and Tian (2005) show that the truncation error in the Britten-Jones and Neuberger (2000) implied variance is negligible if the truncation points are more than two standard deviations from the current underlying price. See also Prokopczuk and Wese Simen (2015).

[4]As a robustness check, we consider a higher sampling frequency of 1-minute and obtain qualitatively similar findings. We do not tabulate these results for brevity.

We also obtain end-of-day S&P 500 index options data for the period starting in January 1996 and ending in August 2015 from IvyDB OptionMetrics. 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. These options are of the European type. Generally, weekly options expire on the Friday of each week, except the third Friday of each month when the standard options expire.[5,6] The first weekly option on the S&P 500 index appears in October 2005. The data available from the Chicago Board of Option Exchange (CBOE) indicate 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 to 50 % towards the end of our sample (Andersen et al., 2017).[7]

Given the maturity structure of S&P 500 index options, it is possible to use weekly options of different expiration dates to construct the daily, weekly and monthly implied variance each month. We analyze the sample period when the weekly option contracts are available for trading. Accordingly, we restrict our attention to the period extending from January 2006 to the end of August 2015. We start the sample in January 2006, rather than October 2005 when the first weekly option contract starts trading, to allow

[5] At the time of writing, the term-structure of weekly options can include 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*.

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

[7] The reader may ask: why do market participants trade short-term options? As pointed out by Andersen et al. (2017), answering this question is complicated by the lack of official data on the identity/profile of market participants who take positions in these contracts. However, the authors show that there is no significant change in the trading activity of these contracts around important scheduled macroeconomic announcements such as those of the monetary policy rate and the monthly employment report. Thus, they conclude that the increased trading activity in these contracts is primarily driven by a desire to improve short-term risk management. As more official data on short-term options emerge, it would be interesting to analyze the profile of key players in this market segment. Doing so would help shed more light on the economic mechanism behind our results. As our main focus is on modeling realized variance, we leave this avenue of research for future work.

for trading interest in the weekly options to improve. Note, however, that starting the sample in October 2005 does not materially affect our findings.

We process the option data as follows. We discard observations with missing or zero prices. We implement this filter separately using bid and ask prices. 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 S&P 500 index and that of the dividend yield, both obtained from OptionMetrics, with the panel of options data.

Our analysis involves monthly observations of all key variables. We retain the last day of each month when we jointly observe the daily, weekly and monthly implied variance. Typically, this observation day is the last Thursday of each month. An upshot of this is that the daily implied variance depends mainly on the weekly options that expire the following trading day. The weekly implied variance relates to the weekly options expiring at the end of the following week. The monthly implied variance depends on the weekly options expiring at the end of the next month. This argument also implies that the interpolation across maturities has a minimal effect on our results since we are as close as possible to the targeted maturities.[8]

By using a monthly observation frequency, we aim to avoid concerns related to overlapping observation biases discussed in Christensen and Prabhala (1998). Panel A of Table 1 presents the summary statistics on the term-structure of implied variance. The term-structure displays a hockey-stick pattern. The daily, weekly and monthly implied variance have average (annualized) values of 0.059, 0.050 and 0.053, respectively. The short-end of the term-structure is more volatile than the long-end. The skewness and kurtosis coefficients indicate that the distribution of the daily implied variance is distinct from that of the other maturities.

Figure 1 plots the time-series of daily, weekly and monthly (annualized) implied volatilities. Several patterns are worth highlighting. First, the series rise and fall together, indicating that they are positively correlated. Second, they rise during bad economic times, such as the 2008 global recession. Third, the short-term series is higher than that of monthly horizon during periods of economic downturn. This observation suggests that short-term implied variance may contain information that is different from that of the monthly implied variance, which has received a lot of attention from extant studies (Jiang and Tian, 2005; Bekaert and Hoerova, 2014). Panel B of Table 1 presents some supportive evidence in this regard. It shows that the correlation between the daily and monthly implied variance series is positive yet modest (0.64).

[8]As a robustness check, we do not interpolate the constant maturity contracts but instead simply use the option contracts with maturity closest to the target maturity. On average, the "real" maturity of the options are 1, 6 and 27 calendar days for the daily, weekly and monthly horizons, respectively. Furthermore, our main conclusions are unchanged. Thus, we conclude that the method of interpolation plays a minimal role in our results. These results are available upon request.

3 HAR–IV

This section starts by establishing the predictive power of the short-term implied variance for next month's realized variance. This motivates the development of the HAR–IV model. We then explore the channels through which the predictability result arises.

3.1 Is Short-Term Implied Variance Informative About Realized Variance?

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

$$RV_{t+1}^{M} = \alpha + \beta IV_{t}^{X} + \epsilon_{t+1} \tag{8}$$

where α is the intercept. β denotes the slope parameter. IV_t^X is the implied variance of time-to-maturity X, where X can be the daily (D), weekly (W) or monthly (M) maturity. ϵ_{t+1} is the residual of the regression at $t+1$.

If implied variance is informative about future realized variance, we expect the slope parameter to be significantly different from zero. Panel A of Table 2 reports positive Newey and West (1987) corrected t-statistics (with 3 lags) of 11.551, 5.061 and 5.899 for the slope parameters in univariate regressions involving the daily, weekly and monthly implied variance, respectively. These findings establish that implied variance of daily, weekly and monthly horizons individually predict realized variance. The positive t-statistics are consistent with the notion that implied variance predicts realized variance with a positive sign. Thus, our results confirm and extend the findings of earlier studies that focus only on the monthly implied variance (Jiang and Tian, 2005; Busch et al., 2011) to the implied variance of shorter maturities.[9]

Note that, of all three maturities, the highest t-statistic relates to the daily implied variance. This suggests that the shortest maturity is a highly significant predictor of realized variance, a conclusion borne out by its predictive power as well ($Adj\ R^2 = 0.743$). In contrast, the monthly implied variance which has been extensively studied in the literature displays a lower predictive power ($Adj\ R^2 = 0.635$).

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 against such an approach. The findings are also difficult to reconcile with the expectations hypothesis. This theory posits that, of all maturities, the monthly implied variance is the best predictor of monthly realized variance. Clearly, the finding that the daily implied variance achieves a higher predictive power than the monthly implied variance challenges this idea.

Multivariate Evidence We now analyze the incremental information content of the daily implied variance relative to the monthly implied variance. To shed light on this question, we include all three maturities in an encompassing model, which we refer to as the HAR–IV model:

$$RV_{t+1}^{M} = \alpha + \beta IV_t^D + \gamma IV_t^W + \phi IV_t^M + \epsilon_{t+1} \tag{9}$$

[9]The earlier literature surveyed in Poon and Granger (2003) tests the unbiasedness hypothesis on the monthly implied variance. As discussed in Chernov (2007), Prokopczuk and Wese Simen (2014) and Kourtis et al. (2016), the unbiasedness hypothesis hinges on the counterfactual assumption of a constant variance risk premium. In light of these studies, we do not test this hypothesis.

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

The penultimate row of Panel A of Table 2 reports that the slope estimate associated with the daily implied variance (0.416) is similar in magnitude to that obtained in the univariate regression model (0.389). Moreover, the parameter estimate remains significant ($t - stat = 4.361$). Taken as a whole, the results indicate that the short-term implied variance contains information that is different from that of the monthly implied variance.

The predictive power of the HAR–IV model ($Adj\ R^2 = 0.859$) is 35 % higher than that afforded by the monthly implied variance alone. Since the monthly implied variance includes the daily implied variance, the reader may find this result surprising. In order to understand how the HAR–IV can markedly improve on the predictive power of the monthly implied variance, it is useful to recall that a univariate regression of the monthly realized variance on a constant and the lagged monthly implied variance implicitly imposes the restriction that all the components of the monthly implied variance, including the daily implied variance, predict realized variance with the same coefficient. Thus, the markedly superior performance of the HAR–IV suggests that this restriction is strongly rejected by the data.

3.2 Dissecting the Predictability

Having established the information content of the short-term implied variance series for realized variance, we now seek to understand 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 return. More formally, we have:

$$RV_{t+1}^{M} = CV_{t+1}^{M} + JV_{t+1}^{M} \tag{10}$$

where RV_{t+1}^{M} is the (annualized) monthly realized variance for the calendar month starting at time t and ending at time $t+1$. CV_{t+1}^{M} and JV_{t+1}^{M} are the monthly continuous and jump variations of the asset returns computed over the month ending at $t+1$, respectively.

This insight suggests that there are two channels through which short-term implied variance may be informative about next month's realized variance. The first possibility is that the short-term implied variance series contain information about the continuous variation of returns. The second possibility is that short-term implied variance series are 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 finite sample properties than the bipower variation. Thus, we use the $MedRV$ estimator of the continuous variation of returns:

$$CV_{t+1}^{M} = \frac{12m\pi}{(6-4\sqrt{3}+\pi)(m-2)} \sum_{j=0}^{N-1} \sum_{k=3}^{m} \text{median}(|r_{t+1-\frac{j}{N},k}|, |r_{t+1-\frac{j}{N},k-1}|, |r_{t+1-\frac{j}{N},k-2}|)^2 \tag{11}$$

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

Using Equations (10) and (11), it is straightforward to extract the jump variation component:[10]

$$JV_{t+1}^{M} = RV_{t+1}^{M} - CV_{t+1}^{M} \tag{12}$$

3.2.b Continuous Variation

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

$$CV_{t+1}^{M} = \alpha + \beta IV_{t}^{X} + \epsilon_{t+1} \tag{13}$$

where all variables are as previously defined.

Panel A of Table 3 shows that each of the three maturities of 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 daily implied variance boasts the highest Newey and West (1987) corrected t-statistic. Relatedly, its predictive power ($Adj\ R^2 = 0.718$) is larger than that of the monthly implied variance ($Adj\ R^2 = 0.651$), confirming our benchmark results. Combining all maturities in the encompassing model yields an $Adj\ R^2$ of 0.850. Moreover, the daily implied variance remains highly significant. This result echoes our earlier conclusion that the information embedded in the short-term implied variance series

[10] 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 variation. The decomposition results are qualitatively similar to those of our benchmark approach. As a result, we do not tabulate these findings. Furthermore, the conclusion is similar, irrespective of whether we use the bipower variation or $MedRV$.

is different from that of the monthly implied variance.

3.2.c Jump Variation

We now estimate the following forecasting regression:

$$JV_{t+1}^{M} = \alpha + \beta IV_{t}^{X} + \epsilon_{t+1} \tag{14}$$

where all variables are as previously defined.

Panel B of Table 3 summarizes the empirical evidence. The results of univariate regressions indicate that each of the three maturities predicts the monthly jump variation. Thus, we are able to extend the work of Busch et al. (2011), who document that the monthly implied variance predicts the jump variation of S&P 500 returns, to shorter maturities of implied variance. The table also reveals that the daily implied variance outperforms the monthly implied variance in terms of predictive ability. Its Newey and West (1987) t-statistic is higher than that of the monthly implied variance. Relatedly, its predictive power ($Adj\ R^2 = 0.742$) is larger than that of the monthly implied variance ($Adj\ R^2 = 0.250$). Including all three maturities in the same regression, we can see that the daily implied variance is the most significant predictor of the jump variation. Furthermore, the predictive power of the encompassing model ($Adj\ R^2 = 0.766$) is very similar to that of the regression that only includes the daily implied variance as a predictor. In light of this, we conclude that short-term options contain valuable information about the future monthly jump variation.

Summarizing, the predictive power of the daily implied variance series is consistent with both the continuous and jump channels. This contrasts with the monthly implied variance, which is mainly informative about the continuous variation of index returns.

4 HAR–IV vs. HAR–RV

The previous section shows that the HAR–IV model, which uses information from the term-structure of implied variance yields a higher predictive power than a model that only includes the monthly implied variance. It is worth noting that the HAR–IV model uses information from the term-structure of implied variance in a manner that is reminiscent of the HAR–RV model of Corsi (2009). The reader may thus wonder: how does the HAR–IV model compare to the HAR–RV model?

Answering this question is important because a large literature surveyed by Poon and Granger (2003) seeks to ascertain whether options data contain superior information compared to historical data. When presented with the results of these studies, it is tempting to compare the sophisticated HAR–RV model to a simple model that relies solely on the monthly implied variance to predict realized variance. Such evidence is relatively difficult to interpret. The HAR–RV model uses information from three maturities of the term-structure of historical variance whereas the alternative model uses only one maturity of the term-structure of implied variance. The difference in the forecasting performance of these two models could therefore be due to (i) differences in the specification of the models and/or (ii) differences in the quality of historical and options data.

4.1 In-Sample Evidence

Following Corsi (2009), we define the HAR–RV model as follows:

$$RV_{t+1}^{M} = \alpha + \beta RV_t^D + \gamma RV_t^W + \phi RV_t^M + \epsilon_{t+1} \tag{15}$$

where α is the intercept. β, γ and ϕ are the slope parameters. RV_t^D and RV_t^W are the (annualized) daily and weekly realized variance at time t, respectively. These series are defined as follows:

$$
\begin{aligned}
RV_t^D &= 252 \times \sum_{k=1}^{m} r_{t,k}^2 && (16) \\
RV_t^W &= 52 \times \sum_{j=0}^{4} \sum_{k=1}^{m} r_{t-\frac{j}{N},k}^2 && (17)
\end{aligned}
$$

where N is the number of trading days in the month that starts at $t-1$ and ends at t. All other variables are as previously defined.

Figure 2 shows the dynamics of the term-structure of historical volatility. The patterns are broadly similar to those of the implied volatility series (see Figure 1), suggesting that the components of the HAR–RV model may be very similar to those of the HAR–IV model. Panel B of Table 1 confirms this intuition. The correlation between the historical variance and the implied variance of the same maturity is generally higher than 0.71. In light of these findings, it is interesting to formally compare the forecasting performance of the HAR–IV and HAR–RV models.

The results of univariate regressions in Panel B of Table 2 show that each of the three maturities of historical variance is a significant predictor of the monthly realized variance. The lagged monthly realized variance displays the highest t-statistic ($t - stat = 11.940$) and an $Adj\ R^2 = 0.538$. It is interesting to note that the predictive power of the lagged monthly realized variance is lower than that of the lagged monthly implied variance ($Adj\ R^2 = 0.635$) presented in Panel A of the same table. This result is consistent with the finding of Jiang and Tian (2005).

Table 2 shows that the HAR–IV model boasts a stronger predictive ability ($Adj\ R^2 =$

0.859) than the HAR–RV model ($Adj\ R^2 = 0.775$). Thus, we conclude that option data do help to achieve superior variance forecasting performance.

Pursuing our analysis, we estimate the nesting model below:

$$RV^M_{t+1} = \alpha + \beta IV^D_t + \gamma IV^W_t + \phi IV^M_t + \psi RV^D_t + \omega RV^W_t + \eta RV^M_t + \epsilon_{t+1} \tag{18}$$

where all variables are as previously defined.

The last row of Panel A of Table 2 shows that all parameters associated with the term-structure of implied variance are statistically significant. Of all maturities of the historical variance, only the loading on the daily component remains statistically significant. The $Adj\ R^2$ of the nesting model is, in relative terms, 20 % higher than that of the HAR–RV model. Using the $Adj\ R^2$ presented in the last rows of Panels A and B of Table 2, we can conduct an F-test to test the null hypothesis that β, γ and ϕ in Equation (18) are jointly equal to zero. Straightforward calculations show that the $F-stat = 70.775$. Clearly, this is above the relevant critical value, indicating that the term-structure of implied variance contains information that is in addition to that of the term-structure of the historical variance.

4.2 Out-of-Sample Evidence

The preceding analysis shows that the HAR–IV model yields more accurate variance forecasts than the HAR–RV model in-sample. We next investigate whether the predictability results extend out-of-sample.

We use a rolling window containing 6 years of data to estimate the forecasting models in Equations (9) and (15).[11] Equipped with the parameter estimates and the latest observations of the forecasting variables, we then generate a forecast for next month's realized variance. If the forecast is higher (lower) than the highest (lowest) monthly realized variance observed in the rolling estimation window, we set the forecast to the highest (lowest) in-sample observation of the monthly realized variance. By imposing this filter, we are able to avoid the situation where the forecast of variance could be negative (Patton and Sheppard, 2015).[12]

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. To do so, we compute the following four loss functions: mean squared percentage error (MSPE), mean absolute percentage error (MAPE), mean squared error (MSE) and mean absolute error (MAE). These loss functions are defined below:

$$MSPE = \frac{1}{T}\sum_{t=1}^{T}\left(\frac{RV_t^M - E_{t-1}(RV_t^M)}{RV_t^M}\right)^2 \tag{19}$$

$$MAPE = \frac{1}{T}\sum_{t=1}^{T}\left|\frac{RV_t^M - E_{t-1}(RV_t^M)}{RV_t^M}\right| \tag{20}$$

$$MSE = \frac{1}{T}\sum_{t=1}^{T}\left(RV_t^M - E_{t-1}(RV_t^M)\right)^2 \tag{21}$$

$$MAE = \frac{1}{T}\sum_{t=1}^{T}\left|RV_t^M - E_{t-1}(RV_t^M)\right| \tag{22}$$

where T is the total number of out-of-sample forecasts. $E_{t-1}(RV_t^M)$ is the expectation at time $t-1$ of the variance to be realized at the end of month t. All other variables are as previously defined.

[11] The rolling window estimation follows Bollerslev et al. (2016), among others. As a robustness check, we consider a recursive window. Our untabulated results lead to very similar conclusions. Given the similarity of the results based on the rolling and recursive schemes, we only focus on the rolling window framework.

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

Table 4 reports the ratio of the loss function [name in row] associated with the model [name in column] over that of the benchmark HAR–RV model. Therefore, an entry equal to 1 indicates that the model [name in column] does as well as the benchmark HAR–RV model. Entries lower than 1 suggest that the model [name in column] achieves lower average forecasting error than the HAR–RV model. Conversely, entries that are greater than 1 indicate that the forecast errors of the model [name in column] are higher than those of the HAR–RV model.

Focusing on the entries reported under the header "HAR–IV", we can see that the figures vary between 0.260 and 0.545. This set of numbers reveals that the HAR–IV model can reduce the forecasting errors of the benchmark HAR–RV by up to 74 %. We formally test the null hypothesis that the difference in average forecast errors is equal to zero. To this end, we implement the Giacomini and White (2006) test and report the corresponding test statistic in Table 5. The test statistic follows a Chi-squared distribution. The entries at the intersection of the row labelled "HAR–IV" and the column headed "HAR–RV" show the chi-squared test statistic of the null hypothesis that the difference between the average forecast errors of the HAR–RV model and that of the HAR–IV model is equal to zero. These figures are statistically significant for nearly all loss functions which is consistent with the evidence of Table 4. We thus conclude that the HAR–IV model achieves significantly smaller forecasting errors than the HAR–RV model.

4.3 Economic Value

Next, we shed light on the economic value of the documented superior predictability. To this end, we evaluate the portfolio choice of an investor with quadratic utility function who invests a fraction of her wealth ω_t in the risky stock market and the rest $(1 - \omega_t)$

in the risk-free asset. To fix the notation, $E_t(r^M_{t+1})$ denotes the expectation at time t of the future monthly return on the risky stock. The riskless asset yields a return rf_t. The risk aversion parameter is γ. $E_t(RV^M_{t+1})$ is the expectation at time t of the variance for the month ending at $t+1$. The optimization problem of the investor can be expressed as follows:[13]

$$\max_{\omega_t} \omega_t E_t(r^M_{t+1}) + (1-\omega_t)rf_t - \frac{\gamma}{2}\omega_t^2 E_t(RV^M_{t+1}) \tag{23}$$

It is straightforward to show that:

$$\omega_t \quad = \quad \frac{E_t(r^M_{t+1}) - rf_t}{\gamma E_t(RV^M_{t+1})} \tag{24}$$

Equation (24) establishes that the optimal share of wealth invested in the risky asset depends on the expected market excess return as well as the expectation of future variance. Since our paper studies the predictability of realized variance, rather than that of excess returns, we follow Bollerslev et al. (2017) and frame the problem in terms of risk modelling. Heeding on their approach, we assume the agent targets a constant ratio of the expected monthly market excess return over the square root of the expected monthly realized variance. With a slight abuse of terminology, we term this ratio SR:[14]

$$SR = \frac{E_t(r^M_{t+1}) - rf_t}{\sqrt{E_t(RV^M_{t+1})}} \tag{25}$$

[13]The optimization problem of an investor with quadratic utility is equivalent to maximizing a linear combination of mean and variance. This is true irrespective of the distribution of asset returns. We refer the interested reader to Campbell and Viceira (2002) for an excellent treatment of this topic.

[14]It is important to emphasize that, strictly speaking, this ratio is not the conditional expectation of the future Sharpe ratio. See also Bollerslev et al. (2017).

Combining Equations (24) and (25), we can see that:

$$\omega_t = \frac{SR}{\gamma\sqrt{E_t(RV_{t+1}^M)}} \tag{26}$$

We then analyze the certainty equivalent rate of return (CE), which is the risk-free rate of return that the investor is willing to accept instead of following the proposed timing strategy:

$$CE = \frac{1}{T}\sum_{t=1}^{T}\left(\frac{SR^2}{\gamma}\times\frac{\sqrt{RV_t^M}}{\sqrt{E_{t-1}(RV_t^M)}}+rf_t-\frac{SR^2}{2\gamma}\times\frac{RV_t^M}{E_{t-1}(RV_t^M)}\right) \tag{27}$$

Similar to the analysis in Section *4.2*, we use a rolling window containing the past 6 years of monthly observations to estimate the variance forecasting models in Equations (9) and (15). We use the parameter estimates together with the latest observations of the forecasting variables to forecast next month's realized variance. Furthermore, we assume a constant $SR = 0.3$ as in Bollerslev et al. (2017).[15] Finally, we consider different values of the risk aversion parameter, namely 3, 4 and 5.

Table 6 reports the difference between the certainty equivalent associated with model [name in column] and that of the HAR–RV model. If the model [name in column] provides more economic value than the HAR–RV model, then its certainty equivalent will be greater than that of the benchmark HAR–RV model. As a result, the entries in the corresponding column should be positive.

The first column indicates whether we allow for leverage in the spirit of Campbell and Thompson (2008). If leverage is restricted, we cap the share of wealth invested in the

[15] We also consider alternative values, i.e. $SR = 0.2$ and $SR = 0.4$, and obtain very similar conclusions. For example, assuming $SR = 0.4$, an investor with a risk aversion parameter equal to 3 is willing to pay 4.399 % per year to switch from a strategy that is based on the HAR–RV model to a strategy based on the HAR–IV model.

risky stock at 1.5 (see Equation (26)). This restriction could matter when comparing the information content of historical data relative to options data. Because option markets are forward-looking, implied variance tends to rise as the economy is about to move from a period of expansion to a period of recession. As a result, the timing strategy will invest less in the stock market ahead of recessions. In contrast, because the historical variance is backward looking, the variance forecast emanating from the HAR–RV is likely to be low, since the historical variance estimates are low before recessions, resulting in a larger allocation to the risky stock. Imposing the leverage restriction enables us to ascertain the extent to which the results may be driven by the aforementioned effect.

The entries in Table 6 show that the HAR–IV model improves on the utility achieved by the timing strategy based on the HAR–RV model. For example, an agent with risk aversion parameter of 3 would pay up to 3.887 % per year to switch from a timing strategy based on the HAR–RV model to a strategy based on the HAR–IV model. This improvement speaks directly to the economic value of using information from the term-structure of implied variance as opposed to the term-structure of historical variance. Qualitatively similar results emerge for other values of the risk aversion parameter. Finally, the leverage restriction has very little effect on the results.

5 What About Measurement Errors?

The summary statistics in Table 1 show that the daily implied variance is more volatile than its monthly counterpart. Furthermore, it is also moderately correlated with the monthly implied variance. One may attribute this set of results to measurement errors in the daily implied variance. If short-term options attract little trading interest, the corresponding option implied variance will be measured with errors. To the extent that the measurement errors are heteroscedastic, they may materially affect the strength of the relationship between the implied variance and the future realized variance. This section analyzes the potential impact of measurement errors on our variance forecasts.

5.1 Framework

Historical Variance The problem of measurement errors in the realized variance series is reminiscent of the error-in-variables problem addressed in a recent study by Bollerslev et al. (2016). The authors study the impact of heteroscedastic measurement errors in the inputs of the HAR–RV model on the model's forecasting performance. They introduce a model that allows for a time-varying sensitivity to the lagged daily realized variance. In our setup, we define the HAR-RQ-D model as:

$$RV_{t+1}^{M} = \alpha + (\beta + \psi ME_t^D)RV_t^D + \gamma RV_t^W + \phi RV_t^M + \epsilon_{t+1} \tag{28}$$

where α is the intercept. β, ψ, γ and ϕ are slope parameters. ME_t^D is the conditional variance of the measurement error at time t associated with the lagged daily variance. All other variables are as previously defined.

They also propose a more flexible model, which we term HAR–RQ–F, that captures the heteroscedasticity of the measurement errors associated with each of the three maturities:

$$RV_{t+1}^{M} = \alpha + (\beta + \psi ME_t^D)RV_t^D + (\gamma + \omega ME_t^W)RV_t^W + (\phi + \eta ME_t^M)RV_t^M + \epsilon_{t+1} \tag{29}$$

where α is the intercept. β, ψ, γ, ω, ϕ and η are slope parameters. ME_t^W and ME_t^M denote the conditional variance at time t of the measurement error associated with the weekly and monthly horizons, respectively. All other variables are as previously defined.

We follow Bollerslev et al. (2016) and express the conditional variance of measurement errors as the square root of the realized quarticity of corresponding maturity (Barndorff-Nielsen, 2002).[16] Despite its simple intuition and theoretical grounding, this modelling approach does not directly apply to the HAR–IV model because there is no analog of the realized quarticity for implied variance. We are thus forced to specify a model for the conditional variance of measurement errors in the implied variance series.

Implied Variance We assume that the conditional variance of the measurement error in the implied variance depends on the illiquidity of the options that underpin its calculations.[17] That is, when the option contracts attract little trading interest, we expect the associated implied variance to be noisy.

We use the inverse of the open interest to proxy for the illiquidity of the option contracts.[18] For each trading day and option maturity, we record the open interest of each OTM option contract that passes the filters discussed in Section *2.2*. Next, we average the open interest figures across all these options, thus obtaining an estimate of the open interest of that specific maturity. We do this for all maturities of the same trading day. We then linearly interpolate the open interest across maturities to obtain the

[16]To investigate whether the results are robust to the functional form specification, we consider the logarithmic (instead of the square root) function in Section *6.2* and reach very similar conclusions.

[17]Our approach is related to the work of Aït-Sahalia and Yu (2009) who links the statistical measurement of microstructure noise in the underlying price to financial measures of stock illiquidity.

[18]We use the inverse of open interest to facilitate the interpretation of the variable as a proxy for illiquidity. When open interest is equal to zero, we assume that the illiquidity is equal to 0. As a robustness check, we discard observations associated with an open interest equal to 0 and obtain similar results. Furthermore, we consider other proxies of illiquidity, i.e. the dollar bid–ask spread and the proportional bid–ask spread, that do not suffer from this issue. Section *6.3* presents qualitatively similar results.

open interest of daily, weekly and monthly horizons.[19] On average, the daily, weekly and monthly open interest amount to 125,522, 504,240 and 1,192,033 contracts, respectively. The standard deviation of the open interest positions is similar to the average level.

We then estimate the HAR–IVME–D model:

$$RV_{t+1}^{M} = \alpha + (\beta + \psi ME_t^D) IV_t^D + \gamma IV_t^W + \phi IV_t^M + \epsilon_{t+1} \tag{30}$$

where α is the intercept. β, ψ, γ and ϕ are the slope parameters. ME_t^D is the conditional variance at time t of the measurement error, that we proxy with the square root of the inverse of the daily open interest. All other variables are as previously defined.

The model in Equation (30) only captures the impact of heteroscedastic measurement errors affecting the daily implied variance. However, one could also account for the heteroscedastic measurement errors linked to the weekly and monthly implied variance. This leads to the HAR–IVME–F model:

$$RV_{t+1}^{M} = \alpha + (\beta + \psi ME_t^D) IV_t^D + (\gamma + \omega ME_t^W) IV_t^W + (\phi + \eta ME_t^M) IV_t^M + \epsilon_{t+1} \tag{31}$$

where ME_t^W and ME_t^M denote the conditional variance at time t of the measurement errors associated with the weekly and monthly implied variance, respectively. We proxy the conditional variance of measurement errors with the square root of the inverse of the corresponding maturity open interest. All other variables are as previously described.

[19] We interpolate across maturities to be consistent with the construction of the constant maturity implied variance series defined in Section *2.1*. As a robustness check, we do not linearly interpolate across maturities and simply select the open interest of the options with maturities that are closest to the daily, weekly and monthly horizons. We obtain very similar results, suggesting that our findings are not driven by the interpolation across maturities. This result is to be expected. As we explain in Section *2.2*, our empirical design suggests a very small role for the interpolation method in our results.

5.2 Results

In-Sample Results The model in Equation (30) enables us to formally evaluate the impact of measurement errors in the daily implied variance on the predictability of realized variance. The second row of Panel A of Table 7 reports a statistically significant loading associated with the interacted daily variance of measurement errors ($t - stat = 9.000$). Thus, we conclude that the relationship between next month's realized variance and the lagged daily implied variance is time-varying. Analyzing the predictive power of the HAR–IVME–D model ($Adj\ R^2 = 0.920$), we observe an improvement relative to the HAR–IV model ($Adj\ R^2 = 0.859$).

The last row of Panel A of Table 7 shows that the slope parameters linked to the weekly and monthly measurement errors are not statistically significant. We conclude that there is little gain to be achieved by accounting for the heteroscedastic measurement errors linked with the weekly and monthly implied variance series.

Out-of-Sample Results Consistent with the evidence of Bollerslev et al. (2016), Table 4 shows that the HAR–RQ–F improves on the performance of the benchmark HAR–RV model.[20] We can also see that the HAR–IVME–D model slightly improves on the out-of-sample performance of the HAR–IV. However, Table 5 reveals that these improvements are not statistically significant. Comparing the HAR–IVME–D and the HAR–RQ–D models, we find that the option-based model leads to significantly more accurate forecasts. Viewed as a whole, the results suggest that accounting for heteroscedastic measurement errors does not materially affect our conclusions on the usefulness of options data.

[20] Note that the HAR–RQ–D model underperforms the HAR–RQ–F model. This result is not surprising given earlier studies. Bollerslev et al. (2016) recommend correcting the historical variance of maturity corresponding to the forecasting horizon. Hence, in our framework, we should account for the heteroscedastic measurement error in the lagged monthly realized variance. This insight helps understand why the HAR–RQ–D model performs so differently relative to the HAR–RQ–F.

Economic Value Table 6 shows the results of the economic value exercise. It reports several noteworthy findings. First, the strategies based on the extensions of the HAR–RV model do not improve utility when compared to the strategy based on the HAR–RV. This result holds for all values of the risk aversion parameter. Second, the HAR–IVME–D and the HAR–IVME–F models lead to higher certainty equivalent rate of returns than the HAR–IV model.

Taken as a whole, the results indicate that the forecasting model based on options data delivers more economic value than the model that uses historical data. This is true, for all values of the risk aversion parameter.

6 Robustness Checks

This section presents several robustness checks. First, we check whether our results are due to the wildcard option feature. Second, we use a logarithmic rather than a square root functional form for the measurement error correction. Third, we consider alternative proxies for the illiquidity of options.

6.1 Wildcard Option Feature

One may worry about the wildcard option feature induced by the asynchronous closing times of the spot and derivatives markets (Harvey and Whaley, 1992). Briefly, the S&P 500 option market closes at 3:15 PM while the spot index market is last updated 15-minute earlier. This asynchronicity between the closing times could affect our results.

Since the 5^{th} of March 2008, OptionMetrics reports the last option quote at 3:00 PM rather than 3:15 PM. We repeat our main analyses using the period from the 5^{th} of March 2008 to the end of August 2015 and obtain very similar results. Table 8 presents the results of the related out-of-sample analysis. They are similar to our benchmark findings. The HAR–IV model achieves lower forecasting errors than the HAR–RV model and its extensions. We thus conclude that the wildcard option feature does not materially affect our conclusions.

6.2 Functional Form

As previously discussed, the square root specification is motivated by the work of Bollerslev et al. (2016) who point out that this functional form is imbued with a certain robustness. Nonetheless, one may wonder about the sensitivity of our main results to the functional form.

To shed light on this, we assume that the conditional variance of the measurement error is simply proxied by the logarithm (rather than square root) of the realized quarticity and the logarithm of the inverse of the open interest. Table 9 presents these results. We can see that the HAR–IV model and its extensions that account for heteroscedastic measurement errors outperform the HAR–RQ–D and the HAR–RQ–F models. This result reinforces our initial finding that options data are more informative about future realized variance than historical data.

6.3 Illiquidity Proxy

The preceding analysis relies on (the inverse of) open interest as a proxy for the illiquidity of option contracts. We now consider alternative proxies based on the dollar and proportional bid–ask spreads. The dollar bid–ask spread is simply the difference between the ask and the bid prices. The proportional bid–ask spread is the dollar bid–ask spread divided by the average of the ask and bid prices.

Our approach in this regard is similar to that described in Section *5.1*. Briefly, we compute these proxies for each option contract that underpins the implied variance calculation. We then average across all these options (of the same maturity), obtaining an illiquidity proxy for each maturity. Finally, we interpolate across maturities to obtain the illiquidity proxies of daily, weekly and monthly horizons.

Panels A and B of Table 10 present results that are qualitatively similar to those of Table 4. The similarity of the results suggests that the exact illiquidity proxy has very little bearing on our main findings.

7 Conclusion

We exploit the recent introduction of weekly options on the S&P 500 index to analyze the information content of short-term options for realized variance. Our results reveal that short-term implied variance predicts next month's realized variance. This result arises because short-term implied variance strongly predicts both the continuous and jump variations of S&P 500 returns. The information content of the short-term implied variance is not subsumed by the monthly implied variance, suggesting that the different maturities do not contain the same information.

We combine the daily, weekly and monthly implied variance to create the HAR–IV model. Empirically, the HAR–IV model outperforms the HAR–RV model and its extensions. This is true both in- and out-of-sample. The superior predictive ability of the HAR–IV model is economically meaningful, as evidenced by utility gains of up to 3.887 % per year relative to the HAR–RV. Our evidence is robust to concerns related to measurement errors.

References

Aït-Sahalia, Y. and Yu, J. (2009). High frequency market microstructure noise estimates and liquidity measures. *Annals of Applied Statistics*, 3(1):422–457.

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. (2015). The risk premia embedded in index options. *Journal of Financial Economics*, 117(3):558–584.

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

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.

Bollerslev, T., Hood, B., Huss, J., and Pedersen, L. H. (2017). Risk everywhere: Modeling and managing volatility. *Duke Working Paper*.

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.

Campbell, J. Y. and Thompson, S. B. (2008). Predicting excess stock returns out of sample: Can anything beat the historical average? *Review of Financial Studies*, 21(4):1509–1531.

Campbell, J. Y. and Viceira, L. M. (2002). *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors*. Oxford University Press.

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.

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

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

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

Giacomini, R. and White, H. (2006). Tests of conditional predictive ability. *Econometrica*, 74(6):1545–1578.

Harvey, C. R. and Whaley, R. E. (1992). Market volatility prediction and the efficiency of the s & p 100 index option market. *Journal of Financial Economics*, 31(1):43–73.

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.

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

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. (2015). Variance risk premia in commodity markets. *AFA Meetings 2015 Paper*.

Figure 1: Term-Structure of Implied Volatility

This figure shows the evolution of the term-structure of (annualized) implied volatility. The vertical axis shows the volatility estimate on an annualized basis. The horizontal axis indicates the date. The blue line relates to the daily implied volatility. The green line depicts the weekly implied volatility. Finally, the red line shows the dynamics of the monthly implied volatility. All data are sampled at the monthly frequency between January 2006 and August 2015.

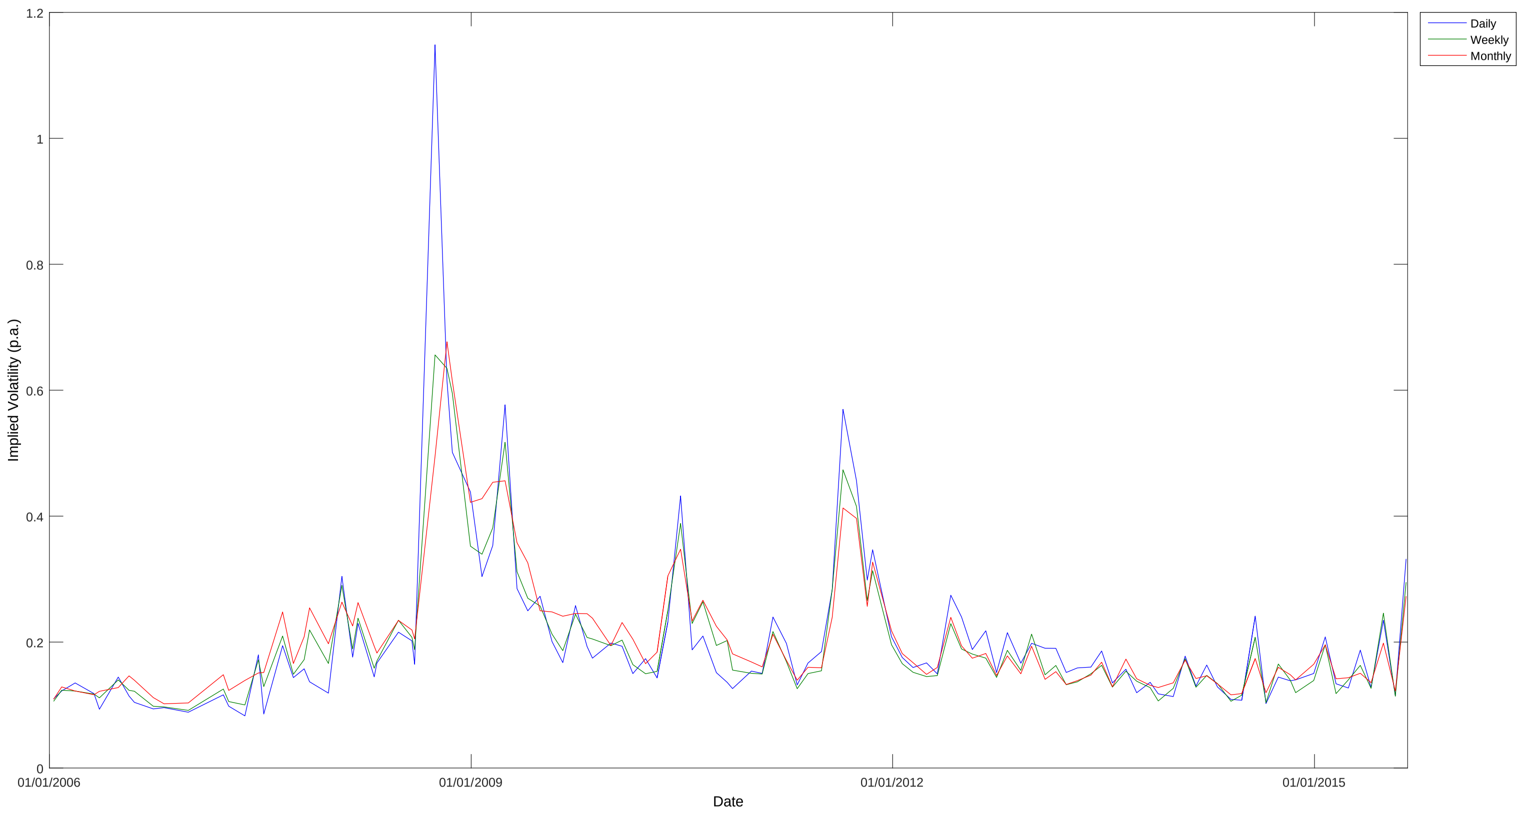

Figure 2: Term-Structure of Historical Volatility

This figure shows the evolution of the term-structure of (annualized) lagged realized volatility. The vertical axis shows the volatility estimate on an annualized basis. The horizontal axis indicates the date. The blue line relates to the lagged daily realized volatility. The green line depicts the weekly lagged realized volatility. Finally, the red line shows the dynamics of the monthly lagged realized volatility. All data are sampled at the monthly frequency between January 2006 and August 2015.

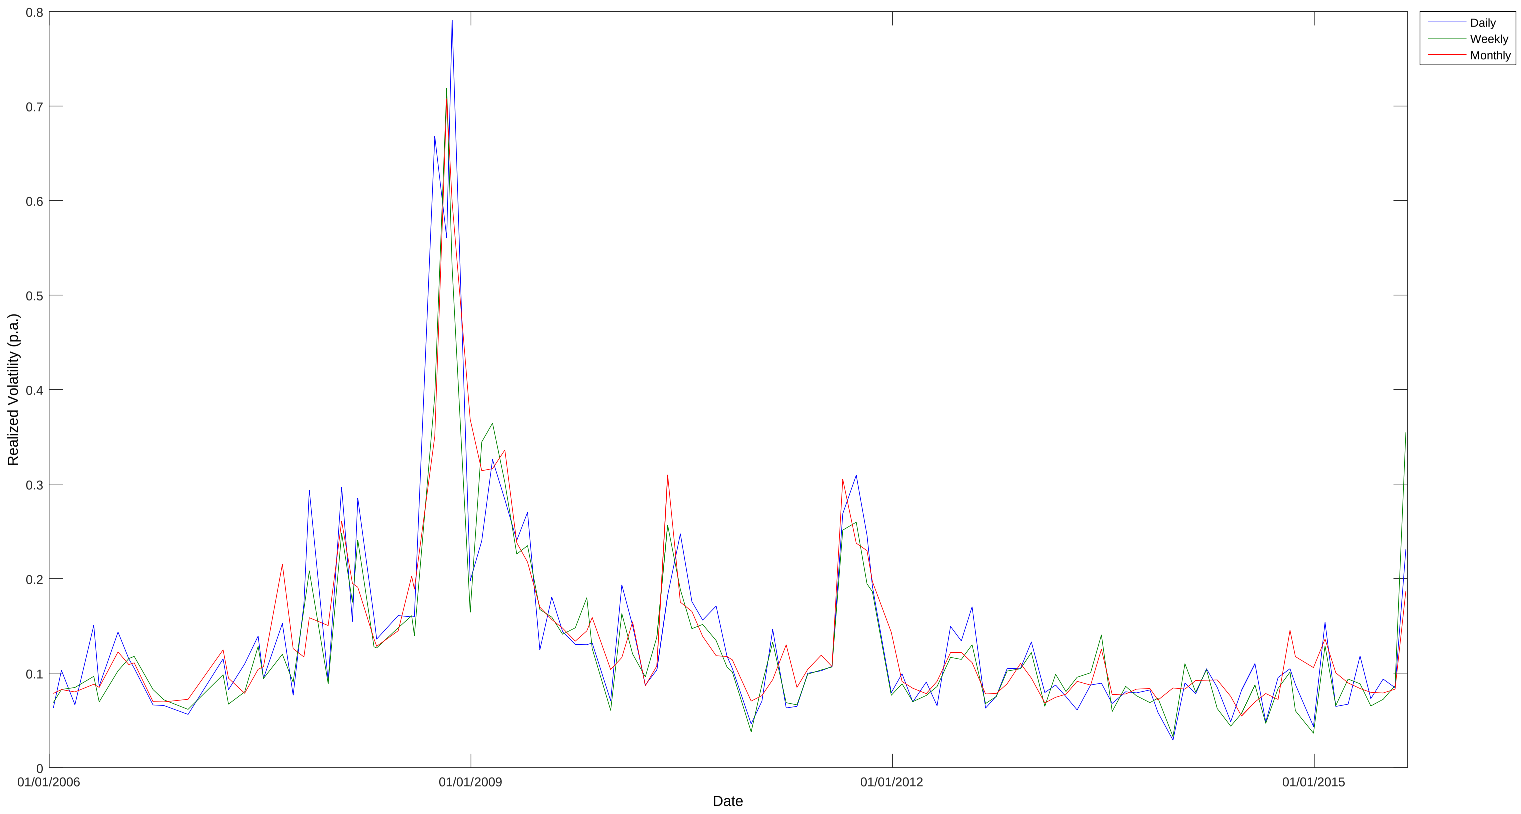

Table 1: Descriptive Statistics

Panel A presents key summary statistics. IV^D, IV^W and IV^M denote the (annualized) implied variance of daily, weekly and monthly horizons, respectively. Similarly, RV^D, RV^W and RV^M denote the (annualized) realized variance of daily, weekly and monthly horizons, respectively. $Mean$ is the average value of the variable [name in row]. Std, $Skew$ and $Kurt$ denote the standard deviation, skewness and kurtosis of the variable [name in row], respectively. Panel B reports the sample correlation between the time-series of the variables [name in row] and [name in column]. The analysis is based on monthly observations. The sample period is from January 2006 to August 2015.

Panel A: Summary Statistics

	$Mean$	Std	$Skew$	$Kurt$
IV^D	0.059	0.136	7.488	67.738
IV^W	0.050	0.071	3.693	17.583
IV^M	0.053	0.067	3.671	18.982
RV^D	0.032	0.078	5.817	39.772
RV^W	0.027	0.059	6.018	45.862
RV^M	0.029	0.061	5.784	40.936

Panel B: Correlation Matrix

	RV^D	RV^W	RV^M	IV^D	IV^W
RV^W	0.78				
RV^M	0.80	0.96			
IV^D	0.71	0.52	0.49		
IV^W	0.86	0.83	0.83	0.84	
IV^M	0.83	0.91	0.93	0.64	0.94

Table 2: Variance Predictability: In-Sample Evidence

This table summarizes the results of regressions of monthly realized variance on a constant and the lagged forecasting variable(s) [name in column]. Panel A mainly focuses on the information content of the term-structure of implied variance. Panel B explores the information content of the term-structure of historical variance. α denotes the intercept parameter. IV^D, IV^W and IV^M denote the (annualized) implied variance series of daily, weekly and monthly horizons, respectively. Similarly, RV^D, RV^W and RV^M denote the (annualized) historical variance series of daily, weekly and monthly horizons, respectively. We present in parentheses the Newey–West corrected t-statistics with 3 lags. *Adj* R^2 is the adjusted R-squared of the regression model. The analysis is based on monthly observations. The sample period is from January 2006 to August 2015.

Panel A: Implied Variance

α	IV^D	IV^W	IV^M	RV^D	RV^W	RV^M	*Adj* R^2
0.006	0.389						0.743
(2.288)	(11.551)						
-0.009		0.753					0.756
(-1.662)		(5.061)					
-0.010			0.732				0.635
(-2.595)			(5.899)				
-0.008	0.416	-0.713	0.899				0.859
(-2.053)	(4.361)	(-1.730)	(4.090)				
0.003	0.415	-0.890	0.529	0.322	0.154	0.131	0.923
(1.494)	(12.213)	(-4.769)	(3.213)	(14.374)	(1.493)	(1.013)	

Panel B: Historical Variance

α	RV^D	RV^W	RV^M	*Adj* R^2
0.007	0.689			0.764
(1.940)	(5.790)			
0.008		0.772		0.555
(2.631)		(9.135)		
0.008			0.741	0.538
(2.643)			(11.940)	
0.006	0.615	0.369	-0.234	0.775
(2.456)	(2.879)	(1.151)	(-0.560)	

Table 3: Continuous vs. Jump Variation

This table dissects the source of the predictability of the monthly realized variance. Panel A presents the results of regressions of the monthly 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 monthly jump variation on a constant and the forecasting variable(s) [name in column]. α denotes the intercept parameter. IV^D, IV^W and IV^M denote the (annualized) implied variance of daily, weekly and monthly horizons, respectively. We present in parentheses the Newey–West corrected t-statistics with 3 lags. $Adj\ R^2$ is the adjusted R-squared of the regression model. The analysis is based on monthly observations. The sample period is from January 2006 to August 2015.

Panel A: Continuous Variation

α	IV^D	IV^W	IV^M	$Adj\ R^2$
0.005	0.360			0.718
(1.955)	(10.532)			
-0.009		0.709		0.759
(-1.861)		(4.928)		
-0.011			0.697	0.651
(-2.820)			(5.689)	
-0.009	0.372	-0.647	0.857	0.850
(-2.284)	(3.954)	(-1.603)	(4.029)	

Panel B: Jump Variation

α	IV^D	IV^W	IV^M	$Adj\ R^2$
0.001	0.029			0.742
(4.221)	(18.002)			
0.001		0.043		0.438
(1.317)		(3.708)		
0.001			0.035	0.250
(1.989)			(3.185)	
0.001	0.045	-0.066	0.042	0.766
(3.452)	(11.947)	(-2.893)	(1.817)	

Table 4: Out-of-Sample Results

This table presents the ratio of the loss function [name in row] of the model [name in column] over that of the HAR–RV model. We use a rolling window of 6 years to estimate the parameters of the forecasting models. HAR–RV is the forecasting model that uses the daily, weekly and monthly lagged realized variance series to predict the monthly realized variance. HAR–RQ–D takes into account the heteroscedasticity of the measurement errors in the lagged daily realized variance. HAR–RQ–F aims to capture the heteroscedasticity of the measurement errors in all three maturities of lagged realized variance. HAR–IV is the forecasting model that simultaneously uses the daily, weekly and monthly implied variance to predict realized variance. HAR–IVME–D accounts for the heteroscedasticity of measurement errors in the daily implied variance. HAR–IVME–F accounts for the heteroscedasticity of measurement errors in all three implied variance components. We proxy the heteroscedasticity of the measurement errors in the realized variance with the square root of the realized quarticity of corresponding maturity. In order to proxy the heteroscedasticity of measurement errors in the implied variance series, we use the square root of the inverse of the open interest of corresponding maturity. The analysis is based on monthly observations.

	HAR–RQ–D	**HAR–RQ–F**	**HAR–IV**	**HAR–IVME–D**	**HAR–IVME–F**
MSPE	2.779	0.902	0.260	0.262	0.315
MAPE	1.819	0.920	0.480	0.468	0.504
MSE	1.553	0.377	0.270	0.236	0.235
MAE	1.632	0.809	0.545	0.507	0.522

Table 5: Out-of-Sample Tests

This table presents the Giacomini and White (2006) test statistic related to the null hypothesis that the difference in the loss function [name in Panel's title] achieved by the model [name in column] and that of the model [name in row] is equal to zero. The test statistic follows a Chi-squared distribution. We use a rolling window of 6 years to estimate the parameters of the forecasting models. HAR–RQ–D takes into account the heteroscedasticity of the measurement errors in the lagged daily realized variance. HAR–RQ–F aims to capture the heteroscedasticity of the measurement errors in all three maturities of lagged realized variance. The HAR–IV simultaneously uses the daily, weekly and monthly implied variance to predict realized variance. HAR–IVME–D accounts for the heteroscedasticity of measurement errors in the daily implied variance. HAR–IVME–F accounts for the heteroscedasticity of measurement errors in all three implied variance components. In order to proxy the heteroscedasticity of measurement errors in the implied variance series, we use the square root of the inverse of the open interest of corresponding maturity. The analysis is based on monthly observations.

Panel A: MSPE

	HAR–RV	**HAR–RQ–D**	**HAR–RQ–F**	**HAR–IV**	**HAR–IVME–D**
HAR–RQ–D	11.811				
HAR–RQ–F	0.118	16.088			
HAR-IV	6.320	15.491	4.250		
HAR–IVME–D	6.442	15.933	4.511	0.005	
HAR–IVME–F	4.962	14.845	3.405	0.748	1.048

Panel B: MAPE

	HAR–RV	**HAR–RQ–D**	**HAR–RQ–F**	**HAR–IV**	**HAR–IVME–D**
HAR–RQ–D	18.244				
HAR–RQ–F	0.392	24.154			
HAR-IV	9.310	21.035	7.444		
HAR–IVME–D	10.230	22.503	9.026	0.143	
HAR–IVME–F	8.473	21.112	7.211	0.280	1.522

Panel C: MSE

	HAR–RV	**HAR–RQ–D**	**HAR–RQ–F**	**HAR–IV**	**HAR–IVME–D**
HAR–RQ–D	16.257				
HAR–RQ–F	1.096	3.826			
HAR-IV	1.325	3.912	1.030		
HAR–IVME–D	1.461	4.154	2.461	0.805	
HAR–IVME–F	1.457	4.134	2.051	0.363	0.003

Panel D: MAE

	HAR–RV	**HAR–RQ–D**	**HAR–RQ–F**	**HAR–IV**	**HAR–IVME–D**
HAR–RQ–D	20.193				
HAR–RQ–F	1.223	17.250			
HAR-IV	4.128	15.496	4.169		
HAR–IVME–D	5.019	16.991	6.747	1.061	
HAR–IVME–F	4.550	16.358	5.872	0.162	0.171

Table 6: The Economic Value of Variance Predictability

The table reports the difference between the certainty equivalent of the timing strategy based on the model [name in column] and that of the strategy based on the HAR–RV model. The first column ("Lev") indicates whether we allow for leverage or not. If we do not allow for leverage (Lev=No), then the share of wealth in the risky asset is capped at 1.5. We consider different values for the risk aversion parameter (γ). We use a rolling window of 6 years to estimate the parameters of the forecasting models. HAR–RV is the forecasting model that uses the daily, weekly and monthly lagged realized variance series to predict the monthly realized variance. HAR–RQ–D takes into account the heteroscedasticity of the measurement errors in the lagged daily realized variance. HAR–RQ–F aims to capture the heteroscedasticity of the measurement errors in all three maturities of lagged realized variance. HAR–IV is the forecasting model that simultaneously uses the daily, weekly and monthly implied variance to predict realized variance. HAR–IVME–D accounts for the heteroscedasticity of measurement errors in the daily implied variance. HAR–IVME–F accounts for the heteroscedasticity of measurement errors in all three implied variance components. We proxy the heteroscedasticity of the measurement errors in the realized variance with the square root of the realized quarticity of corresponding maturity. In order to proxy the heteroscedasticity of measurement errors in the implied variance series, we use the square root of the inverse of the open interest of corresponding maturity. The analysis is based on monthly observations. All values are annualized and expressed in percentage points.

Lev	γ	**HAR–RQ–D**	**HAR–RQ–F**	**HAR–IV**	**HAR–IVME–D**	**HAR–IVME–F**
Yes	$\gamma = 3$	-1.210	-0.362	3.887	4.860	5.078
	$\gamma = 4$	-0.908	-0.272	2.915	3.645	3.808
	$\gamma = 5$	-0.726	-0.217	2.332	2.916	3.047
No	$\gamma = 3$	-1.210	-0.062	4.205	4.834	5.080
	$\gamma = 4$	-0.908	-0.272	2.915	3.645	3.808
	$\gamma = 5$	-0.726	-0.217	2.332	2.916	3.047

Table 7: Variance Predictability: In-Sample Evidence (Measurement Errors)

This table summarizes the results of regressions of monthly realized variance on a constant and the lagged variable(s) [name in column]. Panel A explores the information content of the term-structure of implied variance. Panel B analyzes the information content of the term-structure of historical variance. α denotes the intercept. IV^D, IV^W and IV^M denote the (annualized) implied variance series of daily, weekly and monthly horizons, respectively. RV^D, RV^W and RV^M denote the (annualized) historical variance series of daily, weekly and monthly horizons, respectively. ME^D, ME^W and ME^M are the measurement errors associated with the daily, weekly and monthly components. We proxy the heteroscedasticity of the measurement errors in the realized variance with the square root of the realized quarticity of corresponding maturity. We use the square root of the inverse of the open interest to capture the conditional variance of the measurement errors in implied variance. We present in parentheses the Newey–West corrected t-statistics with 3 lags. *Adj* R^2 is the adjusted R-squared of the regression model. The analysis is based on monthly observations. The sample period is from January 2006 to August 2015.

Panel A: Implied Variance

α	IV^D	IV^W	IV^M	$ME^D \times IV^D$	$ME^W \times IV^W$	$ME^M \times IV^M$	*Adj* R^2
-0.008	0.416	-0.713	0.899				0.859
(-2.053)	(4.361)	(-1.730)	(4.090)				
0.002	0.410	-0.829	0.568	0.032			0.920
(0.810)	(10.624)	(-4.307)	(4.189)	(9.000)			
0.003	0.407	-0.807	0.569	0.032	-0.002	-0.019	0.919
(1.135)	(10.432)	(-4.093)	(4.242)	(8.243)	(-0.591)	(-1.601)	

Panel B: Historical Variance

α	RV^D	RV^W	RV^M	$ME^D \times RV^D$	$ME^W \times RV^W$	$ME^M \times RV^M$	*Adj* R^2
0.006	0.615	0.369	-0.234				0.775
(2.456)	(2.879)	(1.151)	(-0.560)				
0.017	-0.512	0.567	0.043	0.357			0.862
(3.352)	(-1.606)	(2.675)	(0.195)	(3.439)			
0.007	-0.490	0.104	1.026	0.344	0.439	-0.548	0.922
(2.863)	(-2.223)	(0.285)	(4.618)	(6.677)	(3.221)	(-4.898)	

Table 8: Out-of-Sample Results: Wildcard Option Feature

This table sheds light on the impact of the wildcard option. The focus is on the sample period starting from March 05, 2008. The table presents the ratio of the loss function [name in row] achieved by the model [name in column] over that of the HAR–RV model. We use a rolling window of 6 years to estimate the parameters of the forecasting models. HAR–RV is the forecasting model that uses the daily, weekly and monthly lagged realized variance series to predict the monthly realized variance. HAR–RQ–D takes into account the heteroscedasticity of the measurement errors in the lagged daily realized variance. HAR–RQ–F aims to capture the heteroscedasticity of the measurement errors in all three maturities of lagged realized variance. HAR–IV is the forecasting model that simultaneously uses the daily, weekly and monthly implied variance to predict realized variance. HAR–IVME–D accounts for the heteroscedasticity of measurement errors in the daily implied variance. HAR–IVME–F accounts for the heteroscedasticity of measurement errors in all three implied variance components. We proxy the heteroscedasticity of the measurement errors in the realized variance with the square root of the realized quarticity of the corresponding maturity. In order to proxy the heteroscedasticity of measurement errors in the implied variance series, we use the square root of the inverse of the open interest of the corresponding maturity. The analysis is based on monthly observations.

	HAR–RQ–D	**HAR–RQ–F**	**HAR–IV**	**HAR–IVME–D**	**HAR–IVME–F**
MSPE	2.125	0.812	0.298	0.315	0.466
MAPE	1.456	0.818	0.527	0.518	0.617
MSE	1.106	0.256	0.204	0.216	0.237
MAE	1.216	0.682	0.524	0.526	0.571

Table 9: Out-of-Sample Results: Alternative Functional Form

The table presents the ratio of the loss function [name in row] achieved by the model [name in column] over that of the HAR–RV model. We use a rolling window of 6 years to estimate the parameters of the forecasting models. HAR–RV is the forecasting model that uses the daily, weekly and monthly lagged realized variance series to predict the monthly realized variance. HAR–RQ–D takes into account the heteroscedasticity of the measurement errors in the lagged daily realized variance. HAR–RQ–F aims to capture the heteroscedasticity of the measurement errors in all three maturities of lagged realized variance. HAR–IV is the forecasting model that simultaneously uses the daily, weekly and monthly implied variance to predict realized variance. HAR–IVME–D accounts for the heteroscedasticity of measurement errors in the daily implied variance. HAR–IVME–F accounts for the heteroscedasticity of measurement errors in all three implied variance components. We proxy the heteroscedasticity of the measurement errors in the realized variance with the logarithm of the realized quarticity of the corresponding maturity. In order to proxy the heteroscedasticity of measurement errors in the implied variance series, we use the logarithm of the inverse of the open interest of the corresponding maturity. The analysis is based on monthly observations.

	HAR–RQ–D	**HAR–RQ–F**	**HAR–IV**	**HAR–IVME–D**	**HAR–IVME–F**
MSPE	2.757	1.487	0.260	0.177	0.153
MAPE	1.813	1.107	0.480	0.391	0.374
MSE	1.566	0.558	0.270	0.198	0.189
MAE	1.630	0.957	0.545	0.446	0.455

Table 10: Out-of-Sample Results: Alternative Illiquidity Proxies

The table presents the ratio of the loss function [name in row] achieved by the model [name in column] over that of the HAR–RV model. We use a rolling window of 6 years to estimate the parameters of the forecasting models. HAR–RV is the forecasting model that uses the daily, weekly and monthly lagged realized variance series to predict the monthly realized variance. HAR–RQ–D takes into account the heteroscedasticity of the measurement errors in the lagged daily realized variance. HAR–RQ–F aims to capture the heteroscedasticity of the measurement errors in all three maturities of lagged realized variance. HAR–IV is the forecasting model that simultaneously uses the daily, weekly and monthly implied variance to predict realized variance. HAR–IVME–D accounts for the heteroscedasticity of measurement errors in the daily implied variance. HAR–IVME–F accounts for the heteroscedasticity of measurement errors in all three implied variance components. We proxy the heteroscedasticity of the measurement errors in the realized variance with the square root of the realized quarticity of the corresponding maturity. In order to proxy the heteroscedasticity of measurement errors in the implied variance series, we use the square root of the inverse of the dollar bid–ask spread (Panel A) and the proportional bid–ask spread (Panel B) of the corresponding maturity. The analysis is based on monthly observations.

Panel A: Dollar Bid–Ask Spread

	HAR–RQ–D	**HAR–RQ–F**	**HAR–IV**	**HAR–IVME–D**	**HAR–IVME–F**
MSPE	2.779	0.902	0.260	0.234	0.298
MAPE	1.819	0.920	0.480	0.400	0.491
MSE	1.553	0.377	0.270	0.210	0.213
MAE	1.632	0.809	0.545	0.438	0.488

Panel B: Proportional Bid–Ask Spread

	HAR–RQ–D	**HAR–RQ–F**	**HAR–IV**	**HAR–IVME–D**	**HAR–IVME–F**
MSPE	2.779	0.902	0.260	0.240	0.273
MAPE	1.819	0.920	0.480	0.465	0.479
MSE	1.553	0.377	0.270	0.207	0.213
MAE	1.632	0.809	0.545	0.492	0.498