The Informational Content of Implied Volatility in Individual Stocks and the Market*†

Andrey Fradkin

Professor George Tauchen, Faculty Advisor

Duke University

Durham, North Carolina

2008

*The Duke Community Standard was upheld in the completion of this thesis.

†Honors Thesis submitted in partial fulfillment of the requirements for Graduation with Distinction in Economics in Trinity College of Duke University.

Abstract

We examine the informational content of historical and implied measures of variance through an evaluation of forecasts over horizons ranging from 1 to 22 days. These forecasts use heterogeneous autoregressive (HAR) regressions which are constructed with high-frequency data. Our results show that the fit and forecasting ability of models based on historical realized variance (RV) increases with the addition of implied volatility in the regression model. We find that robust regression is better than OLS in forecasting RV outside of the estimation sample. The paper evaluates data from individual equities and the S&P 500. [1]

[1] I would like to thank Professors George Tauchen, Tim Bollerslev, Bjorn Eraker and Victor Todorov for their advice and encouragement. I would also like to thank Tzuo Hann Law for helping me obtain the data for this project. My peers, Peter van Tassel and Warren Davis, in the Honors Workshop provided great support. I would not have been able to write such an in-depth thesis without financial assistance from the Davies Fellowship.

1 Introduction

The volatility of asset returns is of crucial importance in the academic study of financial markets, asset pricing, and risk management. Furthermore, financial instruments based on volatility have become an integral part of the day-to-day business of major corporations. The increased importance of these instruments has coincided with the increased availability of high-frequency data on asset returns. Advancements in asset pricing and volatility modeling have allowed for the information contained in high-frequency data to be used more effectively. Specifically, models based on high-frequency measures of volatility have been found to produce forecasts of volatility that are superior to forecasts done with low-frequency volatility measures. For some of these results see Andersen, Bollerslev, Diebold, and Labys (2003). The implied volatility obtained from option prices provides an alternate way to forecast the future volatility of asset returns.

The implied volatility is an expectation of future volatility under the assumptions of risk neutral preferences and constant volatility. It is often viewed as the market's expectation of future volatility and theoretically incorporates all relevant information including the information contained in high frequency returns.

There is a large literature on the efficiency and general performance of historical forecasts of variance. This literature tests models which use generalized autoregressive conditional heteroskedasticity (GARCH), parametric, non-parametric, and Mincer-Zarnowitz frameworks to forecast variance. Important papers on the subject of volatility forecasts using historical volatility are discussed and thoroughly analyzed in Andersen, Bollerslev, Diebold, and Labys (2003). These papers conclude that simple linear models are as good or better in forecasting than more complex GARCH type models. Andersen, Bollerslev, and Diebold (2006), and Forsberg and Ghysels (2007) test the performance of various forms of these heterogenuous autoregressive (HAR) forecasts. We expand on this literature by being one of the first to test these models on a high-frequency data set which contains individual stocks and the whole market. This paper confirms the earlier finding that realized absolute variation provides the best basis for a HAR forecast and that the effect of jumps, as flagged by the method pioneered by Barndorff-Nielsen and Shephard (2004), on future variance is not significant. Furthermore, we compare a robust regression estimation method to the traditional

OLS estimation method used for these forecasts. The robust method consistently outperforms the OLS method across all models and time periods used.

We empirically explore the extent to which the implied volatility contains information about high-frequency return volatility. This is done by comparing in-sample fit and out-of-sample performance for estimated models using implied volatility, historical variables, and combinations of the two. Furthermore, using high-frequency data available for individual stocks, this paper compares informational content of the implied volatility on stocks to the informational content of implied volatility of the market. The results provide support for the hypothesis that implied forecasts for individual stocks have more informational content than implied forecasts for the market. Furthermore, implied forecasts generally outperform historical forecasts out-of-sample. A combination of implied and historical forecasts is the best of all the models tested, suggesting that historical forecasts and implied forecasts have mutually exclusive information. There are several papers that have explored similar topics. See Jiang and Tian (2007), Fleming (1998), and Andersen, Frederiksen, and Staal (2007). These papers find conflicting results on issues such as the relative informational content of volatility forecasts, implied volatility measures, and historical models. This paper's advantage is that it uses previously untested models and generates results on a larger data sample that includes high-frequency returns for both stocks and the market.

The rest of the paper proceeds as follows. Section 2 contains a discussion of the relevant models of volatility and jumps. Sections 3 and 4 discusses the HAR-RV and Mincer-Zarnowitz classes of models. Section 5 describes the manner in which the realized variance and implied volatility data were obtained and filtered. Section 6 describes the method and justification for robust regressions done in this paper. Section 7 discusses the outcomes from the regressions preformed. Section 8 summarizes the paper and draws attention to the most important results.

2 Model of Volatility and Jumps

In this section we use a model of price movement that incorporates jumps. Consider a log price, $p(t)$, that changes over time as

$$dp(t) = \mu(t)d(t) + \sigma(t)dw(t) \qquad 0 \le t \le T \tag{1}$$

where $\mu(t)dt$ represents the time-varying drift component of the stock. The time-varying volatility of the price movement is represented by $\sigma(t)$ and the $dw(t)$ term is standardized Brownian motion. This is a standard, continuous model of price movements which does not include jumps. Recent literature has suggested that the addition of jumps in the price process is important for theoretical and empirical modeling. The jump processes are added into the following equation.

$$dp(t) = \mu(t)d(t) + \sigma(t)dw(t) + \kappa(t)dq(t), \qquad 0 \le t \le T \tag{2}$$

The non-continuous portion of the price movement is added with the $\kappa(t)dq(t)$ term where q(t) a counting process and $\kappa(t)$ is the magnitude of the jump. There are multiple ways to estimate variation of this process using high frequency financial data. We use realized variance and bipower variation, the two most common and easy to calculate measures of variance. These measures are calculated daily and intraday geometric returns are denoted as, $r_{t,j} = p(t-1+\frac{j}{M}) - p(t-1+\frac{j-1}{M})$, j=1,2,...M , where M is the sampling frequency and t is the day. Throughout this paper the sampling frequency is M=78 which corresponds to 5 minute returns. The first measure of quadratic variation is the Realized Variance

$$RV_t = \sum_{j=1}^{M} r_{t,j}^2 \tag{3}$$

and the alternate measure is the Bipower Variation

$$BV_t = \mu_1^{-2}(\frac{M}{M-1})\sum_{j=2}^{M}|r_{t,j-1}||r_{t,j}| = \frac{\pi}{2}(\frac{M}{M-1})\sum_{j=2}^{M}|r_{t,j-1}||r_{t,j}| \tag{4}$$

where $\mu_a = E(|Z|^a)$, $a > 0$. These two measures were thoroughly investigated in Barndorff-Nielsen and Shephard (2004, 2005) to produce asymptotic results that allow for the separate identification of the continuous and jump components of the quadratic variation. Specifically, they show that as $\frac{1}{M} \to 0$, $BV_{t+1} \to \int_t^{t+1} \sigma^2(t)$ and $RV_{t+1} \to \int_t^{t+1} \sigma^2(s)ds + \sum_{t<s\le t+1} k^2(s)$. Thus, $RV_{t+1} - BV_{t+1} \to \sum_{t<s\le t+1} k^2(s)$.

The difference between the RV and the BV isolates the jump component of the daily volatility. This result can be used to test the hypothesis that a jump occurred on any particular day. The test can be expressed as a z-statistic.

$$z_t = \frac{\frac{RV_t - BV_t}{RV_t}}{\sqrt{(v_{bb} - v_{qq})\frac{1}{M}\max(1, \frac{TP_t}{BV_t^2})}}, \; v_{bb} - v_{qq} = (\frac{\pi}{2})^2 + \pi - 5 \tag{5}$$

$$TP_t = M\mu_{4/3}^{-3}(\frac{M}{M-2})\sum_{j=3}^{M} |r_{t,j-2}|^{4/3}|r_{t,j-1}|^{4/3}|r_{t,j}|^{4/3} \tag{6}$$

The TP is used in this test because, as shown in Barndorff-Nielsen and Shephard (2004), it converges to the integrated quarticity of the price process. That is, $TP_{t+1} \rightarrow \int_t^{t+1} \sigma^4(t)dt$. This provides a scale for the difference between RV and BV. It follows that, $z_t \rightarrow N(0,1)$ as $M \rightarrow \infty$. The test operates under the assumption that there are no jumps. This means that high values of the test suggest the presence of jumps on a particular day. When z_t is sufficiently high then we can reject the hypothesis that there are no jumps. Throughout this paper, a z-statistic at the .999 quartile is used to distinguish a day with jumps from a day without jumps. Huang and Tauchen (2005) use Monte Carlo simulation to demonstrate that the z-statistic shown above is of appropriate size, has good power, and has good jump detection abilities.

3 The HAR Class of Models

This paper relies on a variance forecasting model that forecasts well and is based purely on historical price data. The class of models used in this analysis has become widely used for the purposes of variance forecasting and risk management. These models represent the expectation of future variance if all non-variance data is ignored. Recent literature on forecasting variance has highlighted the fact that simple models often outperform more sophisticated parametric models formally incorporating long-memory processes in out-of-sample forecasts. For example, empirical tests performed in Andersen, Bollerslev, Diebold, and Labys (2003) show that both realized variation vector and univariate autoregression models outperform GARCH type models out-of-sample. These models are linear regressions with past values of variation as independent variables and with current values of variation as dependent variables. This paper uses a variety of heterogeneous autoregressive realized variance (HAR-RV) type models to forecast variance. These models were first developed in Müller et. al (1997) and Corsi (2003) and work by linearly parametrizing the conditional variance of discretely sampled returns. These models have two main advantages over other models. They are very easy to estimate and they capture the long-memory of the variance in a manner which is intuitive.

HAR type models use averaged future RV as the dependent variable and use averages of past values of variance measures as the independent variables. This allows the models to take advantage of information from past price variation. Let the multi-period normalized realized variation over h discrete periods be defined as

$$RV_{t,t+h} = h^{-1}[RV_{t+1} + RV_{t+2} + ... + RV_{t+h}] \tag{7}$$

In this paper, the values 1, 5, and 22 are used for h, referring to daily, weekly, and monthly frequencies respectively. The HAR-RV model can then be expressed as,

$$RV_{t,t+h} = \beta_0 + \beta_D RV_{t-1,t} + \beta_W RV_{t-5,t} + \beta_M RV_{t-22,t} + \varepsilon_{t+1} \tag{8}$$

This variable is generally serially correlated up to at least an order of h-1 and possibly more. In order to obtain heteroskedasticity robust standard errors for the HAR-RV type models this paper uses the Newey-West covariance matrix estimator with a lag of 60 days. The standard HAR-RV model can be expanded to include jumps. The results in Section 2 allow for the separation of the continuous component of the variance from the jump component of the variance. This separation was first introduced in Andersen, Bollerslev, and Diebold (2006) which defined the separate components over a period of time as an average of daily observations where,

$$C_{t+1,\alpha} = I[Z_{TP,t+1} \leq \Phi_\gamma] * [RV_t] + I[Z_{TP,t+1} > \Phi_\gamma] * [BV_t] \tag{9}$$

$$J_{t+1,\alpha} = I[Z_{TP,t+1} > \Phi_\gamma] * [RV_t - BV_t] \tag{10}$$

$$C_{t,t+h} = h^{-1}[C_{t+1} + C_{t+2} + ... + C_{t+h}] \tag{11}$$

$$J_{t,t+h} = h^{-1}[J_{t+1} + J_{t+2} + ... + J_{t+h}] \tag{12}$$

I[.] is the indicator function and Φ_γ is the significance level which is set at .999 as suggested in previous papers. With the continuous and jump component separated, the HAR-RV-CJ model is defined as the regression of RV on the lagged averaged normalized continuous and jump components,

$$RV_{t,t+h} = \beta_0 + \beta_{CD}C_{t-1,t} + \beta_{CW}C_{t-5,t} + \beta_{CM}C_{t-22,t} + \beta_{JD}J_{t-1,t} + \beta_{JW}J_{t-5,t} + \beta_{JM}J_{t-22,t} + \varepsilon_{t,t+h} \tag{13}$$

Andersen, Bollerslev, and Diebold (2006) do not find much persistence in the jump component and they also do not find a large improvement in explanatory power from dividing the continuous and jump components. This paper reaches a similar conclusion from more series of data. Another form of the HAR regression that is used in this paper substitutes the realized absolute value (RAV) as the regressor for the realized variation where RAV is defined as:

$$RAV_t = \mu_1^{-1} M^{-1/2} \sum_{j=1}^{M} |r_{t,j}| \tag{14}$$

$$RAV_{t,t+h} = h^{-1}[RAV_{t+1} + RAV_{t+2} + ... + RAV_{t+h}] \tag{15}$$

$$RV_{t,t+h} = \alpha + \beta_D RAV_{t-1,t} + \beta_W RAV_{t-5,t} + \beta_M RAV_{t-22,t} + \varepsilon_{t+1} \tag{16}$$

The HAR model using RAV (HAR-RV-RAV) was shown empirically to be superior both in-sample and out-of-sample to the HAR-RV and HAR-RV-CJ models on a set of S&P 500 data by Forsberg and Ghysels (2007). Theoretically, the advantage of using RAV is that it is highly robust to jumps and sampling error. Forsberg and Ghysels claim that the asymptotic analysis done by Barndorff-Nielsen, Jacod, and Shephard (2004) shows that jumps do not affect RAV asymptotically and that the sampling error for RAV depends on the second moment whereas the sampling error for RV depends on the fourth moment. It is important to note that RAV has different units than RV. Thus, we also set up a similar model for RAV^2(HAR-RV-RAV^2).

Analogous models are can be used to forecast transformations of RV such as log (RV) and the square root of RV.

$$log(RV_{t,t+h}) = \beta_0 + \beta_D log(RV_{t-1,t}) + \beta_W log(RV_{t-5,t}) + \beta_M log(RV_{t-22,t}) + \varepsilon_{t+1} \tag{17}$$

$$(RV_{t,t+h})^{1/2} = \beta_0 + \beta_D (RV_{t-1,t})^{1/2} + \beta_W (RV_{t-5,t})^{1/2} + \beta_M (RV_{t-22,t})^{1/2} + \varepsilon_{t+1} \tag{18}$$

It is worth noting that the R^2's of these transformed models are not directly comparable to the R^2's of the level RV regressions. That is, the fit of a regression with a dependent variable of RV cannot be directly compared to the fit of a regression with a dependent variable of log(RV). Thus the empirical analysis will examine each of the transformation regressions independently.

4 Mincer-Zarnowitz Regressions

Multiple papers have tried to ascertain whether the implied volatility on a stock option is an unbiased and efficient estimator of future volatility once the risk-premium is taken into account. The standard framework for testing volatility based forecast was first developed in a model by Mincer and Zarnowitz (1969). A regression of the form seen below is usually used. Here α is traditionally considered the bias and β is considered the efficiency. $IV_{t,t+k}$ is the implied volatility at a horizon of k periods forward scaled to a daily level. The horizon of the forecast and the horizon of the implied volatility may differ since implied volatility data are not available for all horizons.

$$RV_{t,t+h} = \alpha + \beta IV_{t,t+k} + \varepsilon_{t+1} \tag{19}$$

A completely efficient and unbiased forecast would have $\alpha = 0$ and $\beta = 1$. However, multiple studies have showed that a standard Black-Scholes at the money implied volatility does not fit these criteria. The reasons for this lack of fit are not considered in this paper, which is only concerned with the forecasting ability of the models. One important study of the relative efficiency of implied volatility is Jiang and Tian (2005). That paper shows that model-free implied volatility is superior to Black-Scholes implied volatility in forecasting future variance. However, it does not use HAR type models for historical forecasting which makes a comparison between implied and historical forecasts difficult. Furthermore, the paper does not test forecasts out-of-sample, which makes its methodology questionable. Another study, Andersen, Frederiksen, and Staal (2007), shows that historical and implied forecasts contain independent information about future variance. This paper responds to these results by including two forms of combined forecasts which lump together historical price variables with implied variables.

$$RV_{t,t+h} = \alpha + \beta IV_{t,t+k} + \beta_{CD} C_{t-1,t} + \beta_{CW} C_{t-5,t} + \beta_{CM} C_{t-22,t} + \beta_{JD} J_{t-1,t} + \beta_{JW} J_{t-5,t} + \beta_{JM} J_{t-22,t} + \varepsilon_{t+1} \quad (20)$$

$$RV_{t,t+h} = \alpha + \beta IV_{t,t+k} + \beta_D RAV_{t-1,t} + \beta_W RAV_{t-5,t} + \beta_M RAV_{t-22,t} + \varepsilon_{t+1} \quad (21)$$

The above regressions are simply combinations of HAR-RV-CJ with IV and HAR-RAV with IV.

5 Data Preparation

The high-frequency data for individual stocks and the SPY index was obtained from the Trade and Quote Database (TAQ) that is available at the Wharton Research Database Service (WRDS). The SPY is an exchange traded fund (ETF) which replicates the performance of the S&P 500. The SPY is traded on the American Stock Exchange (AMEX) and has the same returns as the S&P 500.

The high-frequency data was first formatted into a manageable size by Tzou Hann Law. Law used data for 40 stocks, 10 of which were used in this paper. These 10 stocks were chosen on the basis of the reasonably high open interest on their at-the-money call options throughout the period. The data was checked for anomalies such as non-full trading days which were then removed. Furthermore, it was structured to provide prices at uniform time intervals of 30 seconds. In order to get a more manageable number of observations and to eliminate microstructure noise this paper samples the prices of the securities at 5 minute intervals. The formatted data was then used to construct the RV, RAV, C and J components in accordance with the theory in sections II and III. The time span of the data is from the beginning of 2001 to the end of 2005. The data from 2001-2004 was used to estimate the models while the data from 2005 was used for the forecast evaluation.

Options data was obtained from the OptionMetrics Database accessible on WRDS. Ten equities were chosen for this analysis. For a list of stocks and summary statistics see Table 1. These equities all had an open interest on their options that was high relative to the other stocks that were available for analysis from Law's data. We then selected a unique implied volatility for each day for each stock. This implied volatility was taken directly from OptionMetrics, was at the money, and expired close to a month in the future. This was done so that the implied volatilities reflected the market's expectations of volatility over the next month. The implied volatility from OptionMetrics is approximately equal to the Black-Scholes value of a call for American options. Two different implied volatilities were obtained for the SPY, the S&P 500 index tracker. The first implied volatility was obtained from the model-free VIX volatility index. The other implied volatility, denoted as SPX, was obtained from the implied volatilities on S&P 500 options found on OptionMetrics. The volatilities were then converted into a daily implied variance measure. It is worth noting that implied volatility includes information about overnight volatility whereas the historical measures used in this paper are based solely on the trading day.

6 Robust Regressions

The data used in this paper, like other variance data (see Granger and Poon(2005)) is prone to leverage points and sampling error. Leverage points are individual points which have an extremely large effect on the coefficient estimates of a regression model. Sampling errors may be caused by the absence of non-full trading days which were removed from the sample. These may create small disturbances in the data. The amount of data in the paper is so large that searching for highly influential points manually and trying to find their cause is not likely to yield much insight. Furthermore, the variety of multivariate regressions done in this paper means that some models may be spuriously well fit in-sample and that the distributional assumptions may be wrong. These factors suggest that standard OLS may not be the best way to estimate the parameters. OLS suffers from sensitivity to leverage points and deviations of the data generating process from the model. A sample leverage versus residual plot from a HAR-RV-CJ regression is shown in Figure 1. It is clearly visible that a single point has much more leverage than the others. This point may cause the estimates to be off. Robust regressions offer a way to mitigate the effects of such points. There exists a large literature on robust regressions that is unexplored and untapped by the field of academic finance, specifically as it applies to variance forecasting. Robust estimation methods should seek to accomplish three main goals according to Huber (2004). Firstly, they should have a good efficiency for the assumed model. Secondly, small deviations from the model assumptions should only have a small effect on performance. Thirdly, larger deviations from the model should not completely ruin the model.

This paper seeks to correct for these problems by comparing the performance of a robust regression method based on iteratively reweighed least squared based on M-estimators to the performance of an OLS regression. This paper uses a form of robust regression that is easy to implement in STATA using the "rreg" command. The theoretical justification and explanation of this method is outlined in the Appendix. Further reference on this and similar methods can be found in Huber (2004) and Rousseuw and Leroy (1987). An example of the performance of robust regression versus ordinary least squares is shown in Figure 2. This plot shows the regression of the line $y = x$ where the y variables are augmented with random uniform disturbances and where seven points are modified to be completely off of the trend. In this case the robust regression performs more in accordance with the model. It has a slope of 1.01 and a constant of -.319. On the other hand, the OLS regression has a coefficient of .885 and a constant of 3.01. The robust regression achieves estimates that are closer to the original model. This suggests that robust regression may be better at forecasting variance. This hypothesis is tested in the empirical results section.

7 Empirical Results

The empirical results in this paper are split into three sections. Firstly, we analyze the significance, size, and frequency of jump terms in the individual stock and the market. Secondly, we analyze fit for the in-sample time period of the models and estimation techniques described previously in this paper. Thirdly, we analyze the out-of-sample performance of these models.

7.1 Jumps

Possible jump days can be flagged using the z-statistic shown earlier.The results in sections 2 and 3 are used to separate the jump component of the realized variance from the continuous component on days in which jumps were flagged. The amount of jumps flagged by the z-statistic in individual stocks ranges from 21 to 51. This is a wide variation and suggests that some types of stocks may be more prone to jumps than others. For example Bristol-Meyer Squibb, a pharmaceutical company, has the most jumps. This may be because the stocks of pharmaceutical companies are significantly influenced by information about the successes and failures of specific clinical trials. This information, when announced, could be the cause of the excess jumps.

The average number of jumps for the stocks and for the market is close suggesting that there are no significant differences in the frequency of jumps in markets and individual stocks (See Table 1). On days in which jumps were flagged, jumps comprised 35% of the total daily variance. We examine the significance of the jump component in the HAR-RV-CJ regressions for each of the individual stocks. The findings for the individual stocks confirm the analysis in Andersen, Bollserslev, Deibold (2006) that the jump components are not statistically significant and slightly increase the R^2 of the regression. Furthermore, the jump coefficients are often of opposite signs. This further suggests that the jump terms are not useful in this regression because there is no obvious reason that one jump term term would be positive and another negative. These results are shown in Tables 2, 3, and 4 for monthly, weekly, and daily regressions. The significance results are based on Newey-West errors with a lag of 60 days in order to control for autocorrelation and heteroskedasticity in the error term. Similar results were obtained for the log and square root forms of these regressions. It is worth noting that the constant term in these regressions is always significant, suggesting that the HAR-RV-CJ regressions are biased.

7.2 In-Sample Results

The informational content of different variance measures is assessed by comparing the success of forecasts. The in-sample results are important for this evaluation because they produce easily comparable adjusted R^2's for each model. Furthermore, the fit of the models for individual stocks can be directly compared to the fit of the stocks for the index as a whole. Multiple models are estimated based on historical data, implied data, and a combination of the two. These models are estimated for three horizons, 1 day, 5 days, and 22 days. The fit is best for the 5 days ahead forecast and lowest for the 1 day ahead forecast. This confirms the intuition that forecasting the smoothed variables averaged over 5 or 22 days is easier than forecasts daily realized variation. The models are estimated using the robust regression technique outlined earlier as well as through standard OLS. The fit for the market and the average of the stocks for the 22 day horizon is shown in Table 5. Additional information can be found in the technical appendix. More detailed values for individual stocks can be found in the technical appendix. There are several interesting findings in the in-sample regressions. Firstly, combined regressions using both implied and historical variables have the best fit. This suggests that there is unique information in both the implied and historical variables. Secondly, when just the historical models are compared, the HAR-RAV and HAR-RAV^2 models outperform the RV based models. This confirms Forsberg and Ghysels (2007) finding that the realized absolute value is a better forecaster of future realized variance then realized variance itself. Additionally, the HAR-RV-CJ model increases the fit by at most .03 showing that the addition of the jump components does not yield much benefit even in-sample.

There is no clear indication that the fit of implied volatility is better than that of the HAR-RAV model. The average improvement from the historical measure to the combined measure is smaller for the market than for individual stocks (See Table 5 and Figures 4-5). Furthermore, the increase in fit from from the implied model to the combined model is greater for the market than for the individual stocks. This suggests that the implied volatility from options is a better predictor for individual stocks than for the market. Closer analysis shows that historical measures are better for some of the stocks while implied measures are better for others. Equivalently, this result states that the implied volatility contributes more to combined forecasts of equity variance than it does to combined forecasts of market variance. For forecasts 22 days in advance, the magnitude of the increase in R^2 from implied volatility to the combined forecast is less than the increase from the historical variance to the combined forecast if the HAR-RV-CJ model is used. However, this result vanishes when the HAR-RV-RAV is used in the combined forecast. This suggests that the implied volatility contains more information than the RV-CJ model but less than the RV-RAV model. Furthermore, when forecasts of time periods 5 days and 1 day in advance are used, the implied volatility becomes much worse. This is expected because the implied volatility was constructed to be an expectation over a month long period rather than over smaller periods.

7.3 Out-of-Sample Results

In this section we report the results from the out-of-sample prediction of realized variance calculated from the previously discussed regressions. It is worth noting that the evaluation period, 2005, is only a year long. Thus, there are at most 12 independent month-long periods to test the forecast. Furthermore, 2005 was a remarkable year for its low volatility and even more importantly its low variance of volatility. This suggests that the mean squared errors (MSE) are probably lower than if another, more volatile, period was used to evaluate the forecasts. However, the fact that the variance is low in 2005 does not pose a serious problem for the following analysis. There is no reason to think that the performance of robust regression compared to OLS would change in a higher volatility period.

There are several important empirical relationships which we investigate in the out-of-sample analysis. We first look at the difference in performance between robust and OLS regressions. The mean squared errors for both estimation periods are reported for all three horizons in tables 6 through 8. These tables also report the ratio of the mean squared error of the robust regression to the mean squared error of the OLS regression. Values for which the ratio is less than one signify that the robust regression performed better. There is a striking pattern across all of the time horizons and models that the robust regressions have smaller mean squared errors than the OLS regressions. The ratios of robust mean squared error to OLS mean squared error are sometimes as low as .260 even for the best performing models. For example, the robust HAR-RAV forecast for Nokia over a 22 day period produced a mean squared error that was about a quarter of the mean squared error produced by the OLS forecast of the same model. This suggests that previous studies using OLS estimates may have significantly underestimated the forecastability of variance. Furthermore, they may have evaluated the informational content in these forecasts inaccurately by not using the information in the best manner. The disparity between robust and OLS estimates is smallest for the 1 day ahead forecasts. This is expected because daily realized variances are not smoothed and are thus the least predictable. Considering the superior performance of the robust estimation method, we base our conclusions about informational content mainly on the robust results. Preliminary work suggests that other robust estimation methods such as quantile regression produced similar increases in forecasting accuracy.

For the 22 day horizon forecasts, the implied volatility Mincer-Zarnowitz regression usually outperforms all the other non-combination forecasts. While the pattern does not hold for all of the stocks analyzed, it holds for most of them. This suggests that the implied volatility has more information about future variance than historical measures. The implied volatility is even a good predictor of the future values of the market indices. It is worth noting that the model-free, VIX, implied volatility has a slightly lower mean squared error than the Black-Scholes implied volatility. This is an indicator that the VIX has more information about future market volatility than the implied volatility from at-the-money SPX options. The results also show that the mean squared error on market forecasts is much smaller than the mean squared error on individual stock forecasts. This suggests that the variance of the market is generally much more predictable than the variance of individual stocks.

One of the curious findings from these forecasts is that the estimates of some of the historical models have a large constant. This constant causes a bias which is seen most clearly for the HAR-RV-CJ model estimated for Citigroup in Figures 6 and 7. The forecast overestimates the realized variance out-of-sample by a large amount. The robust regression mitigates this effect but does not eliminate it. The problems occurs because the in-sample estimation fits the early part of the sample using a constant at the expense of fit during the later part of the sample. The model then overestimates during the later part of the sample and this extends into the forecast. This lack of fit towards the end of the sample continues into the forecast. This problem occurs to a much lesser extent in the HAR-RV-CJ model for the SPY as show in in Figures 8 and 9. The distributions of the forecast errors are the most centered and densely concentrated for implied volatility forecasts. Two abnormally large forecast errors for the implied volatility model occur when there is a jump which the implied volatility predicts but overshoots. After the jump, the implied volatility adjusts to prior levels.

Most of the results from the in-sample analysis hold up out-of-sample. Adding jumps offers little, if any, benefit in forecasting. RAV based measures outperform the realized variance based measures as independent variables in these regressions. For the 22 day forecast horizon, implied volatility based forecasts often have lower mean squared errors than combined forecasts. This further bolsters the case that implied measures generally have more information about future variance than historical measures.

In order to address the low volatility in the 2005, out-of-sample results were obtained for another period. The same models were estimated using the years 2001 and 2002. Mean squared forecast errors were then calculated for the year 2003. These errors are presented in table 9. The most striking result is that the superiority of the robust regressions over OLS remains. Furthermore, it still remains much easier to forecast the variance of the market than it is to forecast the variance of individual equities. These results provide a strong indicator that the conclusions from the full sample are valid over more volatile periods.

8 Conclusion

This paper explores the relative behavior of a set of variance forecasting models in order to better understand the informational content of implied volatility. We find that the magnitude and occurrence of jumps in price does not have significant effects on future variance. This suggests that it is not necessary to separate the continuous and jump component of volatility for the purpose of forecasting. Furthermore, it supports the intuition that discontinuous movements in price are idiosyncratic and thus do not affect the future volatility of returns. Models based on RAV perform the best in fit and in forecasting of all the models using past price data. We find that robust estimation methods outperform OLS in forecasting in almost all cases. The decrease in MSE from using the robust regression over the OLS is often more than fifty percent. This suggests that both HAR and Mincer-Zarnowitz forecasting models suffer from an extreme sensitivity to leverage points.

The estimation of OLS regressions shows that combined models using both historical and implied measures have the best fit. In most cases, the implied volatility provides a better forecast of future volatility over a month long horizon than historical measures. These results suggest that the implied volatility contains unique information about future price movements. The increase in fit between implied and historical models also shows that the implied volatility does not contain all of the information found in historical price data. Furthermore, an analysis of the changes in fit between models indicate that the implied volatility contains more information about individual stocks than it does for the market. We hypothesize that this is because there are possibly predictable idiosyncratic shocks to equities, such as drug trial results or law suits, whereas there are no comparable shocks to the S&P 500 as a whole. A comparison of out-of-sample forecasts using the VIX and the implied volatility of at-the-money calls shows that the VIX is better at forecasting.

These results provide strong evidence that the implied volatility is a good predictor of future volatility. Even though we cannot make claims about efficiency, we can safely say that implied volatility contains important information about future volatility that is not found in historical variance measures. Furthermore, the results show that future variance in individual equity is less determined by past variance than future variance in the market.

9 References

Andersen, T. G., Frederiksen P.H., & Staal, A. D. (2007). The information content of realized volatility forecasts. Preliminary Manuscript.

Andersen, T. G., Bollerslev T., Diebold F. X., & P. Labys (2003). Modeling and forecasting realized volatility. *Econometrica* 71, 579-625.

Andersen, T. G., Bollerslev, T., & Diebold, F. X. (2006). Roughing it Up: Including Jump

Components in the Measurement, Modeling and Forecasting of Return Volatility. *The Review of Economics and Statistics*, 89(4): 701-720.

Bandi, F. & Perron, B.(2001). Long Memory and the Relation Between Implied and Realized Volatility. Working paper. University of Chicago.

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

Barndorff-Nielsen, O., & Shephard, N. (2004b). Measuring the Impact of Jumps on Multivariate Price Processes Using Bipower Variation. Discussion Paper, Nuffeld College, Oxford University.

Barndorff-Nielsen, O., J. Jacod, & Shephard, N. (2004). Limit Theorems for Bipower Variation in Financial Econometrics. Discussion paper, Aarhus and Oxford University.

Corsi, F. (2003). A simple long memory model of realized volatility. Unpublished manuscript. University of Southern Switzerland.

Fleming, J. (1998). The Quality of Market Volatility Forecasts Implied by S&P 100 Index Option Prices. *Journal of Empirical Finance*, 5, 317-345.

Forsberg, L., & Ghysels, E. (2007). Why do absolute returns predict volatility so well? *Journal of Financial Econometrics* 5, 31-67.

Granger, C.W.J., & Poon, S.-H.(2005). Practical issues in forecasting volatility. *Financial Analysts Journal* 61, pp. 45-56.

Huang, X., & Tauchen, G. (2005). The Relative Contribution of Jumps to Total Price Variation. *Journal of Financial Econometrics* 3, 456-499.

Huber, P.J. (2004). *Robust Statistics.* New York: Wiley.

Jiang, G., & Tian, Y. (2005). The Model-Free Implied Volatility and Its Information Content. *Review of Financial Studies*, 18, 1305-1342.

Mincer, J., & Zarnowitz, V. (1969), The evaluation of economic forecasts. In J. Mincer (Ed.). *Economic Forecasts and Expectations.* NBER, New York.

Müller, U. A., Dacorogna, M. M., Davé, R. D., Olsen, R. B., Pictet, O. V., and von Weizsäcker, J. E. (1997), Volatilities of Different Time Resolutions - Analyzing the Dynamics of Market Components. *Journal of Empirical Finance*, 4, 213-239.

Rousseeuw, P. J., & Leroy, A. M. (1987). *Robust Regression and Outlier Detection.* New York: Wiley.

10 Appendix

Consider a multivariate regression with unknown parameters $\theta_1, ..., \theta_p$ that are estimated from observations $y_1, ..., y_n$ through the equation

$$y_i = \sum_{j=1}^{p} x_{ij}\theta_j + u_i \tag{22}$$

where x_{ij} are observed coefficients and u_i is the independently distributed error term. Traditionally, the coefficient estimates of θ are obtained by minimizing the sum of squares

$$\min_{\theta} \sum_i (y_i - \sum x_{ij}\theta_j)^2 \tag{23}$$

In order to correct for this problem this paper uses robust regression. The version used in this paper is implemented as 'rreg' in STATA. This regression weighs separate observations using an interated procedure. First, it obtains an OLS estimate for the equation. It then drops points with a very high leverage as indicated by a value of Cook's distance[2] greater than unity. Afterwards, it iteratively computes weights based on absolute residuals using Huber weighing where the minimization problem becomes

$$\min_{\theta} \sum_{i} \rho(\frac{y_i - \sum x_{ij}\theta_j}{\sigma})\sigma \tag{24}$$

In the above equation, σ and θ are iteratively recalculated estimate from the residual, $r_i = y_i - \sum x_{ij}\theta_j$, and the minimum is characterized by the solution to the system of equations below:

$$\sum_{i} \psi(\frac{r_i}{\sigma}) = \sum_{i} w_i r_i x_{ij} = 0,\ w_i = \frac{\psi(\frac{r_i}{\sigma})}{\frac{r_i}{\sigma}} \tag{25}$$

$$\sum_{i} \chi(\frac{r_i}{\sigma}) \quad = \quad 0,\ \chi(x) = x\psi(x) - \rho(x) \tag{26}$$

These equations are then solved using a recursive procedure. First trial values of the θ and σ are selected. Then, the scale step σ^m can be defined as follows:

$$(\sigma^{m+1})^2 = \frac{1}{na} \sum \chi(\frac{r_i}{\sigma^m})(\sigma^m)^2 \tag{27}$$

This can be considered the ordinary variance estimate calculated from metrically Winsorized residuals in the case of the Huber weighing function. The location step can then be written using the solution for τ in the below equation.

$$X^T W X \tau = X^T W \mathbf{r}, \tag{28}$$

$$\theta^{m+1} = \theta^m + \tau \tag{29}$$

[2]Cook's Distance is defined as:

$$D = \sum \frac{(\hat{y}_i - \hat{y}_{j(i)})}{p \cdot MSE}$$

where p is the number of parameters, $\hat{y}_i$is the fitted value of the observation i and $\hat{y}_{j(i)}$is the fitted value of the observation i for a regression which is estimated without the observation i.

W is the diagonal matrix with the diagonal elements w_i calculated from the previous values of the location and scale. The Huber weight functions are defined below.

$$\rho(x) = \begin{cases} \frac{x^2}{2} & |x| < c \\ \frac{c(2x-c)}{2} & |x| \geq c \end{cases} \tag{30}$$

$$\rho^{'}(x) = \psi(x) = \begin{cases} x & |x| < c \\ c & x \geq c \\ -c & x \leq -c \end{cases} \tag{31}$$

These are run until the maximum change in weights is below a specific threshold. Afterwords, another set of iterations is done using Tukey's bisquare weighting which takes care of the previous function's problems with severe outliers. For bisquare weighting the relevant equation is:

$$\psi(x) = x[1 - (\frac{x}{R})^2]^2 \tag{32}$$

Bisquare weights often have trouble converging and may lead to multiple solutions. This is one of the reasons that a Huber estimate for the equation is needed first. The iterative process stops after the difference in bisquare weights is below a threshold. The constants c=1.345 and R=4.685 used by STATA produce an estimate which is about 95% as efficient as the OLS estimates.

11 Tables

Table 1: Abbreviations and Summary Statistics

Ticker	BMY	C	GE	GS	HD	KO
Actual Name	Bristol-Meyer Squibb	Citigroup	General Electric	Goldman Sachs	Home Depot	Coca Cola
Number of Jumps in Data	51	21	24	25	34	32
Mean RV	3.17	2.98	2.65	2.54	3.26	1.69
Ticker	MDT	MOT	NOK	TXN	SPY	SPX
Actual Name	Medtronic	Motorola	Nokia	Texas Instruments	SPY with VIX	SPY with Imp. Vol.
Number of Jumps in Data	47	38	33	45	37	37
Mean RV	2.7	8.89	4.98	8.35	1.3	1.3

Table 2: Coefficient Estimates and Significance of Jump Terms in HAR-RV-CJ Regression - 22 Days Ahead

	BMY	C	GE	GS	HD	KO	MDT	MOT	NOK	TXN	SPY
β_{CD}	0.075**	0.148***	0.121***	0.153***	0.172***	0.184**	0.108**	0.081*	0.156***	0.152***	0.148***
β_{CW}	0.449***	0.125	0.256*	0.257**	0.266**	0.171*	0.338**	0.313***	0.398***	0.439***	0.328***
β_{CM}	0.236	0.348***	0.277**	0.407***	0.191**	0.333***	0.214	0.260**	0.242**	0.112	0.212***
β_{JD}	0.117	0.378**	0.237	-0.064***	0.013	0.185**	-0.094**	0.034	0.052	0.15	-0.135***
β_{JW}	-1.284	-0.052	1.472	-0.148	3.444*	0.876	0.509*	-0.507	-0.590*	-0.631	-0.424**
β_{JM}	8.317	1.388	5.596	-1.300***	3.114	-1.945**	2.380***	2.72	-0.684	4.908	0.046
β_0	0.500*	1.065**	0.797**	0.454**	0.996**	0.532**	0.633***	2.187**	0.821**	1.732**	0.326**

* $p<0.05$, ** $p<0.01$, *** $p<0.001$

Table 3: Coefficient Estimates and Significance of Jump Terms in HAR-RV-CJ Regression - 5 Days Ahead

	BMY	C	GE	GS	HD	KO	MDT	MOT	NOK	TXN	SPY
β_{CD}	0.097***	0.296***	0.190***	0.257***	0.245***	0.318**	0.132**	0.104*	0.192**	0.186***	0.208**
β_{CW}	0.499***	0.161*	0.317***	0.259***	0.455***	0.325**	0.553***	0.501***	0.468***	0.544***	0.447***
β_{CM}	0.286**	0.326***	0.299***	0.423***	0.107	0.192*	0.154	0.217*	0.263**	0.15	0.211**
β_{JD}	0.315**	0.884*	-0.09	-0.143***	-0.039	0.329**	-0.204***	0.015	0.166***	0.336	-0.177**
β_{JW}	-0.271	-0.008	2.994*	-0.244**	0.671	0.719	0.423	0.493	-0.684*	-0.606	-0.617***
β_{JM}	2.491	-0.299	3.539	-0.678*	6.846	-0.669	1.413***	0.034	-0.78	1.83	-0.167
β_0	0.289*	0.634*	0.416*	0.189	0.422**	0.279**	0.314***	1.235**	0.365*	0.802**	0.159**

* $p<0.05$, ** $p<0.01$, *** $p<0.001$

Table 4: Coefficient Estimates and Significance of Jump Terms in HAR-RV-CJ Regression - 1 Day Ahead

	BMY	C	GE	GS	HD	KO	MDT	MOT	NOK	TXN	SPY
β_{CD}	0.155***	0.576***	0.326***	0.550***	0.352***	0.349*	0.116	0.082	0.352***	0.394***	0.329
β_{CW}	0.485***	0.084	0.307***	0.096	0.447***	0.448***	0.627***	0.579***	0.359*	0.355***	0.400*
β_{CM}	0.278**	0.221***	0.244***	0.339**	0.086	0.105	0.15	0.217*	0.250*	0.181*	0.202*
β_{JD}	0.575*	0.600**	1.583	-0.024	-0.821***	-0.062	-0.372***	0.443**	0.387***	0.616	-0.249
β_{JW}	0.305	1.351*	-0.965	-0.482***	1.308	1.289***	0.278	-0.051	-0.415	-0.091	-0.541**
β_{JM}	0.84	-0.392	4.057	-0.374	4.76	-0.584*	1.210***	0.119	-0.793	0.852	-0.287
β_0	0.213*	0.330*	0.259*	0.078	0.238*	0.167**	0.224***	0.857**	0.209*	0.463**	0.097**

* $p<0.05$, ** $p<0.01$, *** $p<0.001$

Table 5: In-Sample Adjusted R^2 for OLS Regressions 22 Days Ahead

OLS - 22 Days Ahead	SPY	SPX	AVG
HAR-RV	0.45	0.45	0.51
HAR-RV-CJ	0.47	0.47	0.52
HAR-RAV	0.56	0.56	0.58
HAR-RAV^2	0.48	0.48	0.52
Implied Volatility	0.47	0.46	0.56
HAR-RV-CJ + IV	0.52	0.51	0.59
HAR-RAV + IV	0.57	0.57	0.61
% From RV to CJ	0.02	0.02	0.01
% From RAV to CJ	-0.01	-0.01	0.00
% From CJ to Combo	0.05	0.04	0.07
% From IV to Combo	0.05	0.05	0.03
% From RAV to Combo2	0.01	0.01	0.03
% From IV to Combo2	0.10	0.11	0.05

AVG is the average value over all 10 individual stocks

% Signifies Percentage Change in Adjusted R^2 from one model to another

Combo is a model where the regressors from the HAR-RV-CJ Model are combined with IV

Combo2 is a model where the regressors from the HAR-RAV Model are combined with IV

Bold values signify the highest adjusted R^2 in the non-combo models

Table 6: Out-of-Sample Mean Squared Errors for OLS and Robust Regressions 22 Days Ahead

OLS Regression	BMY	C	GE	GS	HD	KO	MDT	MOT	NOK	TXN	SPY	SPX
HAR-RV	0.510	1.330	0.891	0.289	0.919	0.262	0.360	5.886	1.028	3.841	0.110	0.110
HAR-RV-CJ	0.419	1.316	0.879	0.273	0.783	0.252	0.364	5.776	0.909	4.232	0.107	0.107
HAR-RAV	0.338	0.190	0.168	0.448	0.551	0.096	**0.195**	**1.242**	0.366	1.748	0.054	0.054
HAR-RAV^2	0.699	0.991	0.933	**0.262**	0.926	0.274	0.411	4.140	0.913	3.563	0.118	0.118
Implied Volatility	**0.288**	**0.102**	**0.092**	0.288	**0.260**	**0.073**	0.315	3.620	**0.357**	**1.046**	**0.023**	**0.030**
HAR-RV-CJ + IV	0.235	0.063	0.118	0.241	0.311	0.060	0.216	2.798	0.287	1.722	0.036	0.042
HAR-RAV + IV	0.250	0.091	0.078	0.331	0.379	0.058	0.186	1.026	0.310	1.395	0.041	0.043
Robust Regression												
HAR-RV	0.262	0.184	0.134	0.201	0.214	0.113	0.150	2.202	0.220	1.399	0.041	0.041
HAR-RV-CJ	0.257	0.175	0.139	0.197	0.207	0.109	0.155	2.257	0.205	1.478	0.041	0.041
HAR-RAV	**0.165**	0.071	0.059	0.232	0.202	0.051	**0.106**	**0.656**	**0.095**	0.938	0.027	0.027
HAR-RAV2	0.330	0.142	0.139	**0.190**	0.215	0.112	0.139	1.636	0.217	1.285	0.044	0.044
Implied Volatility	0.237	**0.037**	**0.043**	0.212	**0.142**	**0.046**	0.219	1.843	0.305	**0.753**	**0.016**	**0.018**
HAR-RV-CJ + IV	0.145	0.043	0.053	0.207	0.157	0.040	0.120	1.719	0.122	0.976	0.020	0.021
HAR-RAV + IV	0.141	0.055	0.036	0.214	0.182	0.037	0.101	0.591	0.070	0.855	0.019	0.019
$\frac{Robust}{OLS}$												
HAR-RV	0.514	0.138	0.151	0.696	0.233	0.433	0.416	0.374	0.214	0.364	0.375	0.375
HAR-RV-CJ	0.613	0.133	0.159	0.720	0.264	0.430	0.426	0.391	0.226	0.349	0.383	0.383
HAR-RAV	0.489	0.375	0.352	0.517	0.366	0.533	0.542	0.528	0.260	0.537	0.491	0.491
HAR-RAV2	0.472	0.143	0.149	0.724	0.232	0.408	0.337	0.395	0.238	0.361	0.368	0.368
Implied Volatility	0.823	0.364	0.473	0.734	0.545	0.629	0.695	0.509	0.854	0.720	0.706	0.600
HAR-RV-CJ + IV	0.614	0.674	0.453	0.859	0.506	0.655	0.555	0.614	0.425	0.567	0.572	0.495
HAR-RAV + IV	0.565	0.605	0.464	0.648	0.481	0.633	0.542	0.576	0.227	0.613	0.463	0.433

Bold values denote the lowest MSE of any non-combination model

Table 7: Out-of-Sample Mean Squared Errors for OLS and Robust Regressions 5 Days Ahead

	BMY	C	GE	GS	HD	KO	MDT	MOT	NOK	TXN	SPY	SPX
HAR-RV	0.217	0.529	0.296	0.268	0.461	0.095	0.253	2.710	0.301	1.407	0.044	0.044
HAR-RV-CJ	**0.205**	0.533	0.291	0.271	0.400	0.094	0.258	2.712	0.242	1.448	0.042	0.042
HAR-RAV	0.427	0.594	0.244	0.504	0.802	0.116	0.289	2.544	1.181	3.227	0.047	0.047
HAR-RAV^2	0.331	0.314	0.328	**0.248**	0.441	0.107	**0.245**	1.597	**0.233**	**1.280**	0.048	0.048
Implied Volatility	0.236	**0.307**	0.**123**	0.463	**0.281**	**0.048**	0.276	**1.306**	0.339	1.860	**0.022**	**0.018**
HAR-RV-CJ + IV	0.141	0.178	0.061	0.320	0.255	0.030	0.166	1.056	0.083	0.889	0.017	0.017
HAR-RAV + IV	0.320	0.504	0.222	0.429	0.540	0.095	0.281	1.860	1.020	2.763	0.043	0.042
Robust Regression												
HAR-RV	0.131	0.089	0.056	0.246	0.280	0.047	0.198	1.583	0.136	0.912	0.027	0.027
HAR-RV-CJ	0.132	0.086	0.060	0.245	0.277	0.046	0.198	1.554	0.129	0.918	0.030	0.030
HAR-RAV	0.212	0.108	0.069	0.272	0.309	**0.039**	0.163	1.256	0.316	1.499	0.020	0.020
HAR-RAV2	**0.129**	0.070	0.054	**0.229**	0.258	0.049	**0.158**	1.267	**0.110**	0.865	0.032	0.032
Implied Volatility	0.178	**0.065**	**0.038**	0.299	**0.167**	0.048	0.207	**1.079**	0.156	**0.816**	**0.017**	**0.019**
HAR-RV-CJ + IV	0.111	0.051	0.034	0.230	0.228	0.030	0.152	1.024	0.061	0.638	0.016	0.017
HAR-RAV + IV	0.140	0.096	0.057	0.248	0.271	0.033	0.142	1.104	0.267	1.342	0.017	0.018
$\frac{Robust}{OLS}$												
HAR-RV	0.601	0.169	0.189	0.915	0.608	0.494	0.782	0.584	0.453	0.648	0.611	0.611
HAR-RV-CJ	0.646	0.162	0.207	0.904	0.692	0.488	0.769	0.573	0.534	0.634	0.711	0.711
HAR-RAV	0.497	0.183	0.281	0.540	0.385	0.337	0.564	0.494	0.267	0.465	0.434	0.434
HAR-RAV2	0.390	0.224	0.164	0.922	0.585	0.460	0.648	0.793	0.474	0.676	0.655	0.655
Implied Volatility	0.756	0.212	0.308	0.646	0.595	1.021	0.751	0.826	0.460	0.439	0.787	1.031
HAR-RV-CJ + IV	0.782	0.286	0.550	0.720	0.894	1.006	0.914	0.970	0.737	0.718	0.983	0.980
HAR-RAV + IV	0.438	0.191	0.258	0.577	0.502	0.342	0.506	0.594	0.262	0.486	0.399	0.420

Bold values denote the lowest MSE of any non-combination model

Table 8: out-of-sample Mean Squared Errors for OLS and Robust Regressions 1 Day Ahead

	BMY	C	GE	GS	HD	KO	MDT	MOT	NOK	TXN	SPY	SPX
HAR-RV	0.338	0.218	0.175	0.513	0.656	0.100	0.577	3.278	0.253	1.535	0.059	0.059
HAR-RV-CJ	**0.312**	0.148	**0.105**	0.558	0.655	**0.093**	0.593	2.815	**0.208**	1.500	**0.052**	**0.052**
HAR-RAV	0.680	1.363	0.484	0.807	1.232	0.218	0.677	4.942	1.792	5.026	0.092	0.092
HAR-RAV^2	0.417	**0.143**	0.202	**0.483**	**0.626**	0.109	**0.556**	**2.485**	0.220	**1.466**	0.059	0.059
Implied Volatility	13.718	13.366	13.139	15.269	15.309	3.409	6.062	91.126	98.531	94.381	3.497	2.314
HAR-RV-CJ + IV	1.807	2.437	1.886	2.996	2.513	0.624	1.801	16.075	17.558	13.352	0.748	0.438
HAR-RAV + IV	0.583	1.227	0.501	0.773	0.988	0.206	0.681	4.128	1.592	4.516	0.091	0.088
Robust Regression												
HAR-RV	0.294	0.124	0.098	0.496	0.565	0.083	0.536	2.614	0.184	**1.311**	0.052	0.052
HAR-RV-CJ	**0.293**	**0.113**	**0.087**	0.544	0.634	0.087	0.578	2.591	0.186	1.411	0.049	0.049
HAR-RAV	0.395	0.140	0.114	0.513	0.595	0.101	0.529	2.790	0.386	2.242	0.049	0.049
HAR-RAV^2	0.288	0.113	0.107	0.477	**0.538**	**0.080**	**0.496**	**2.321**	**0.177**	1.353	**0.047**	**0.047**
Implied Volatility	7.897	9.070	10.832	11.102	10.462	2.476	4.139	67.224	88.152	81.552	2.480	1.590
HAR-RV-CJ + IV	0.908	0.887	0.766	2.366	1.706	0.179	1.025	12.004	12.360	9.808	0.475	0.235
HAR-RAV + IV	0.344	0.146	0.097	0.505	0.562	0.095	0.511	2.694	0.337	2.085	0.049	0.048
$\frac{Robust}{OLS}$												
HAR-RV	0.868	0.570	0.559	0.967	0.862	0.829	0.929	0.798	0.728	0.854	0.885	0.885
HAR-RV-CJ	0.938	0.762	0.822	0.976	0.967	0.931	0.973	0.921	0.897	0.940	0.930	0.930
HAR-RAV	0.581	0.103	0.235	0.636	0.483	0.464	0.782	0.565	0.215	0.446	0.532	0.532
HAR-RAV^2	0.690	0.787	0.531	0.988	0.860	0.737	0.892	0.934	0.803	0.923	0.800	0.800
Implied Volatility	0.576	0.679	0.824	0.727	0.683	0.726	0.683	0.738	0.895	0.864	0.709	0.687
HAR-RV-CJ + IV	0.502	0.364	0.406	0.790	0.679	0.287	0.569	0.747	0.704	0.735	0.635	0.536
HAR-RAV + IV	0.591	0.119	0.194	0.653	0.568	0.461	0.750	0.653	0.212	0.462	0.531	0.546

Bold values denote the lowest MSE of any non-combination model

Table 9: Test of Previous Results for 22 Days Ahead for the year 2003												
OLS Regression	BMY	C	GE	GS	HD	KO	MDT	MOT	NOK	TXN	SPY	SPX
HAR-RV	2.307	8.029	7.260	1.263	5.590	0.864	1.963	49.261	9.723	39.034	0.339	0.339
HAR-RV-CJ	6.062	8.477	7.258	1.063	5.221	0.901	1.679	48.181	9.756	36.752	0.347	0.347
HAR-RAV	**2.049**	**3.865**	5.062	0.595	4.128	0.684	1.695	**35.938**	**4.942**	**26.572**	**0.221**	**0.221**
HAR-RAV^2	2.680	7.309	6.816	1.159	5.597	0.922	2.525	45.894	9.353	39.319	0.316	0.316
Implied Volatility	2.517	5.665	**3.542**	**0.557**	**3.372**	**0.564**	**0.795**	37.057	4.945	30.439	0.269	0.343
HAR-RV-CJ + IV	3.921	6.333	4.724	0.490	4.060	0.587	0.861	47.567	7.590	32.638	0.278	0.308
HAR-RAV + IV	2.310	4.996	4.239	0.404	3.883	0.556	0.714	45.291	4.485	27.353	0.215	0.227
Robust Regression												
HAR-RV	1.199	1.812	2.580	0.884	1.009	0.273	0.967	26.707	3.565	18.816	0.101	0.101
HAR-RV-CJ	2.102	1.859	2.503	0.739	1.223	0.288	0.662	**21.287**	3.502	18.805	0.093	0.093
HAR-RAV	**1.078**	**1.355**	2.331	**0.538**	**0.974**	0.243	0.922	21.677	2.699	**13.904**	**0.074**	**0.074**
HAR-RAV2	1.402	1.702	2.641	0.798	1.164	0.264	1.385	25.205	3.627	19.394	0.086	0.086
Implied Volatility	2.046	2.981	**1.615**	0.556	1.185	**0.235**	**0.550**	22.364	**2.411**	19.770	0.080	0.099
HAR-RV-CJ + IV	2.028	2.558	1.686	0.339	1.227	0.186	0.373	26.869	2.950	19.248	0.084	0.069
HAR-RAV + IV	1.677	1.838	1.922	0.369	0.935	0.182	0.333	28.209	1.969	15.905	0.065	0.063
$\frac{Robust}{OLS}$												
HAR-RV	0.520	0.226	0.355	0.700	0.180	0.315	0.315	0.493	0.542	0.367	0.482	0.298
HAR-RV-CJ	0.347	0.219	0.345	0.695	0.234	0.320	0.320	0.394	0.442	0.359	0.512	0.267
HAR-RAV	0.526	0.351	0.460	0.903	0.236	0.355	0.355	0.544	0.603	0.546	0.523	0.334
HAR-RAV2	0.523	0.233	0.387	0.689	0.208	0.286	0.286	0.548	0.549	0.388	0.493	0.271
Implied Volatility	0.813	0.526	0.456	0.998	0.351	0.416	0.416	0.691	0.604	0.487	0.649	0.298
HAR-RV-CJ + IV	0.517	0.404	0.357	0.690	0.302	0.317	0.317	0.434	0.565	0.389	0.590	0.301
HAR-RAV + IV	0.726	0.368	0.453	0.914	0.241	0.328	0.328	0.466	0.623	0.439	0.581	0.301
Bold values denote the lowest MSE of any non-combination model												

12 Figures

Figure 1: Leverage Versus Residual Plot for a Sample HAR-RV-CJ Regression

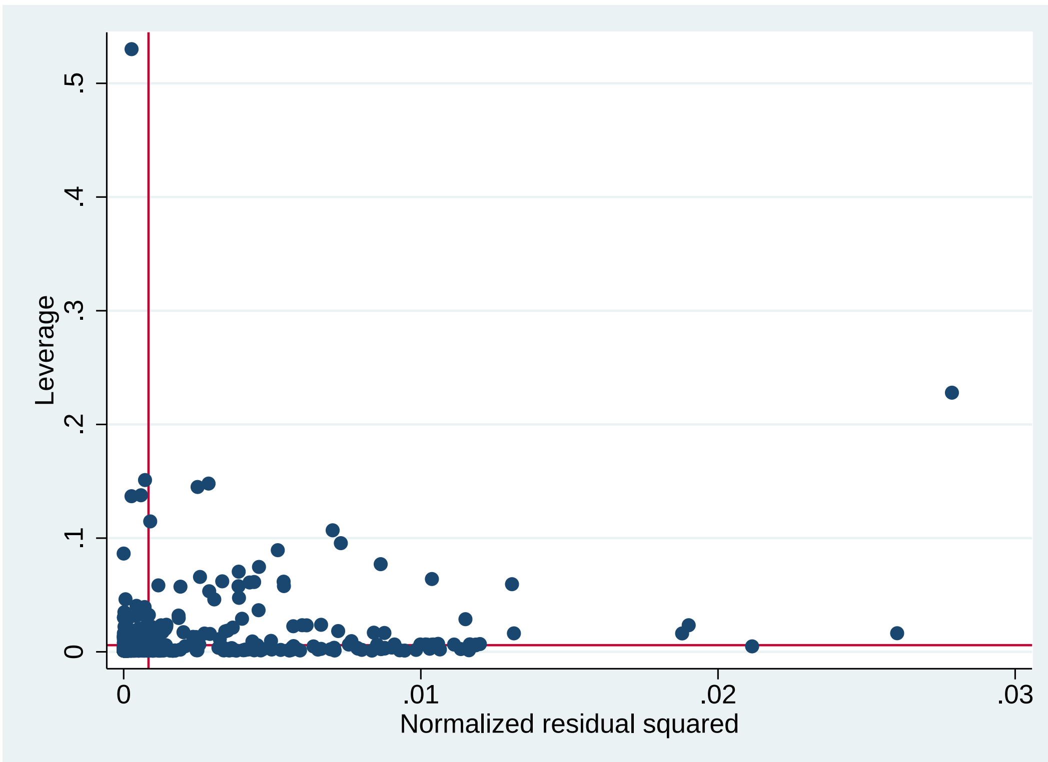

Figure 2: Robust and OLS forecasts of Illustrative Data Set

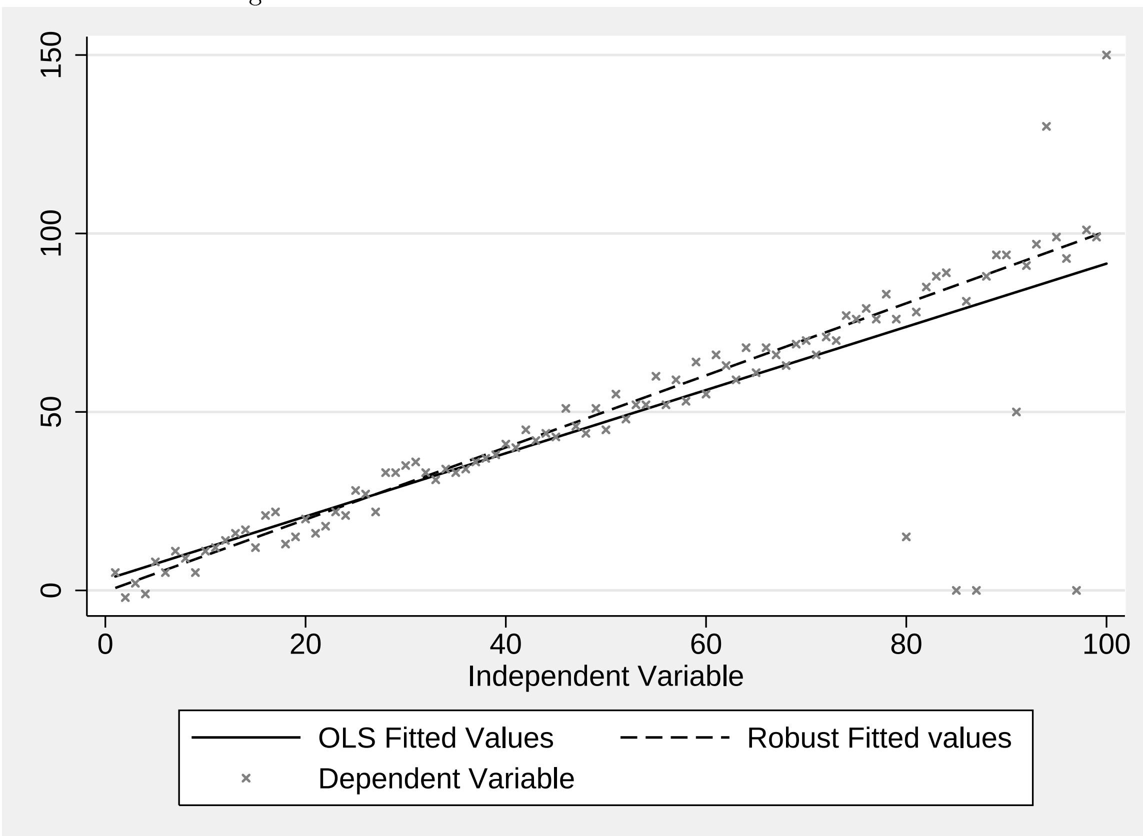

Figure 3: Plots of Functions Used In Robust Regression

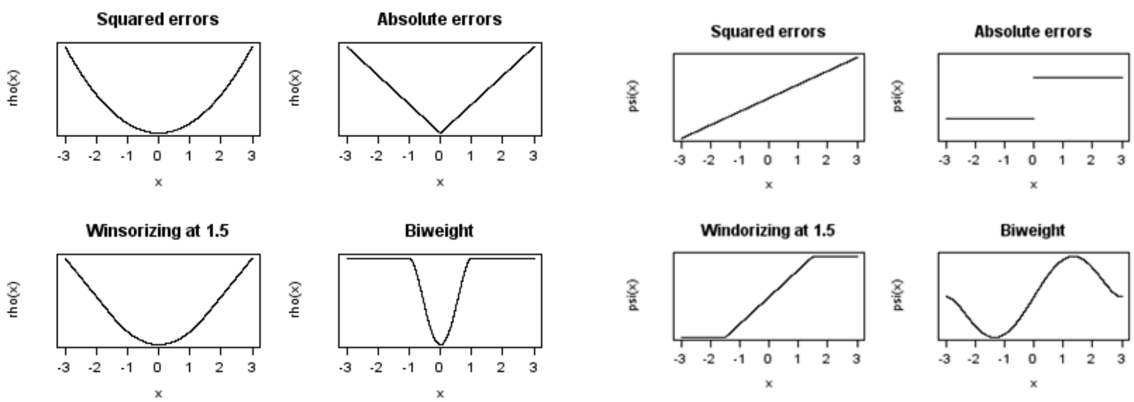

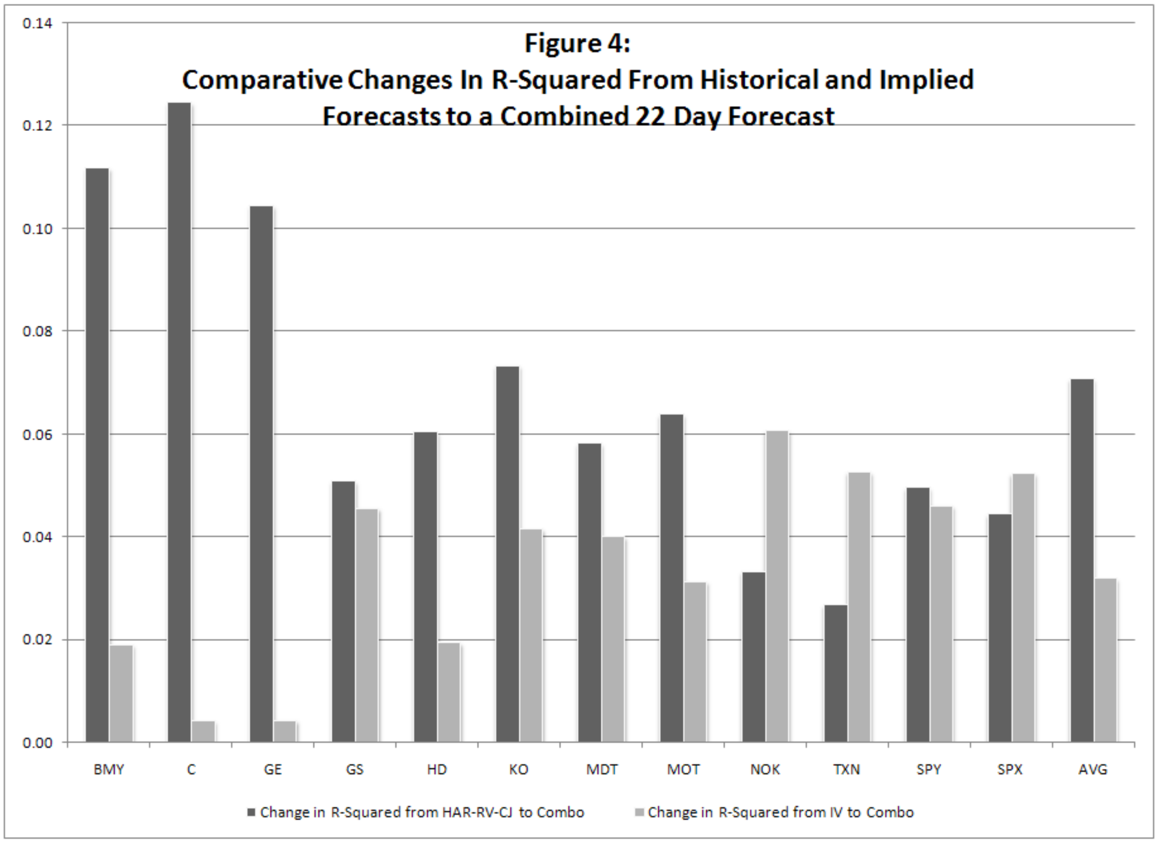
Figure 4:
Comparative Changes In R-Squared From Historical and Implied Forecasts to a Combined 22 Day Forecast
0.14
0.12
0.10
0.08
0.06
0.04
0.02
0.00
BMY
C
GE
GS
HD
KO
MDT
MOT
NOK
TXN
SPY
SPX
AVG
Change in R-Squared from HAR-RV-CJ to Combo
Change in R-Squared from IV to Combo

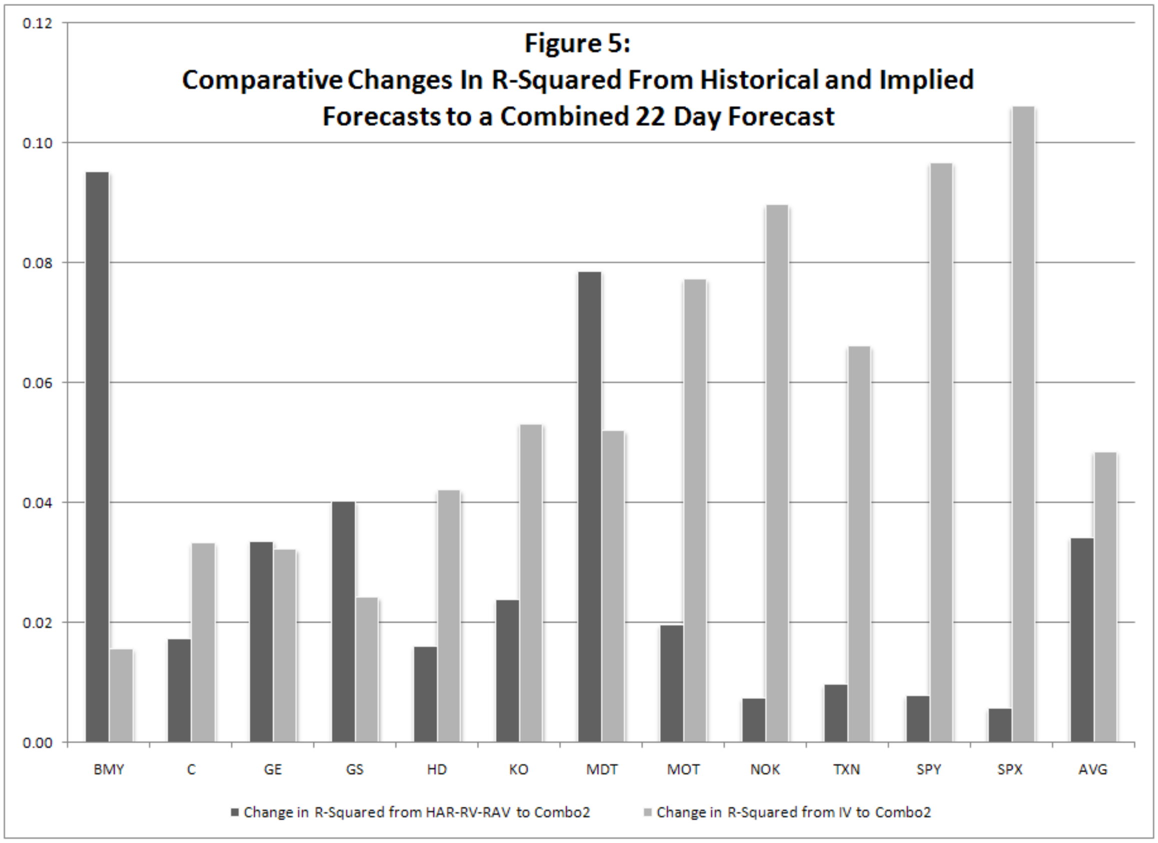
Figure 5:
Comparative Changes In R-Squared From Historical and Implied Forecasts to a Combined 22 Day Forecast
0.12
0.10
0.08
0.06
0.04
0.02
0.00
BMY
C
GE
GS
HD
KO
MDT
MOT
NOK
TXN
SPY
SPX
AVG
Change in R-Squared from HAR-RV-RAV to Combo2
Change in R-Squared from IV to Combo2

Figure 6: Plot of Actual 22 Day Variance Against Different Forecasts for Citigroup Stock

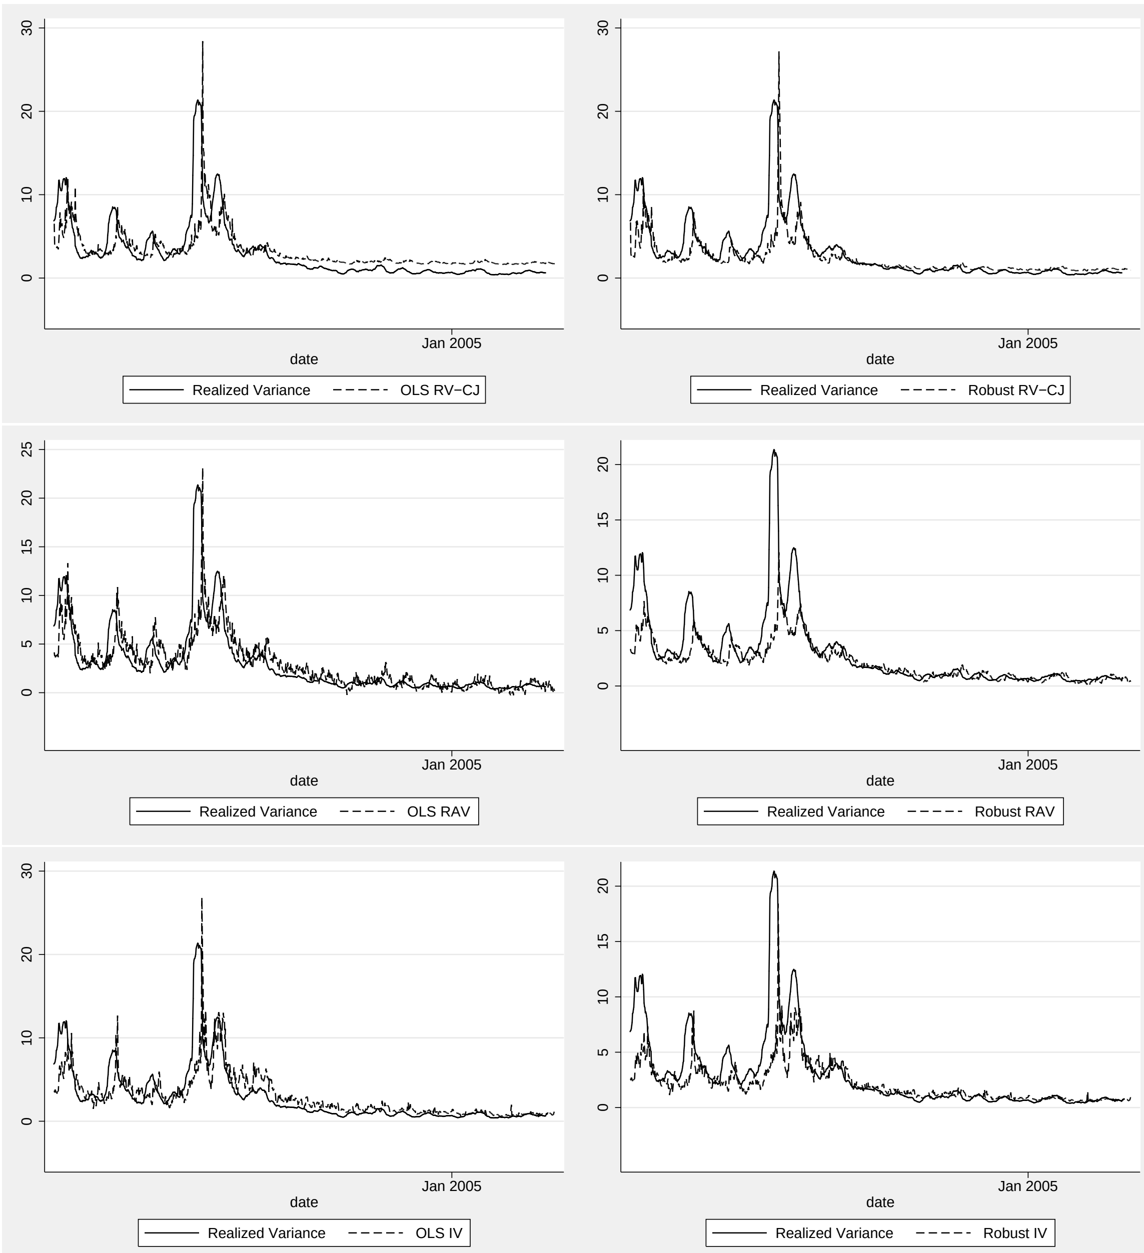

Figure 7: Histograms of Forecast Errors for 22 Day Forecasts of the Variance of Citigroup Stock

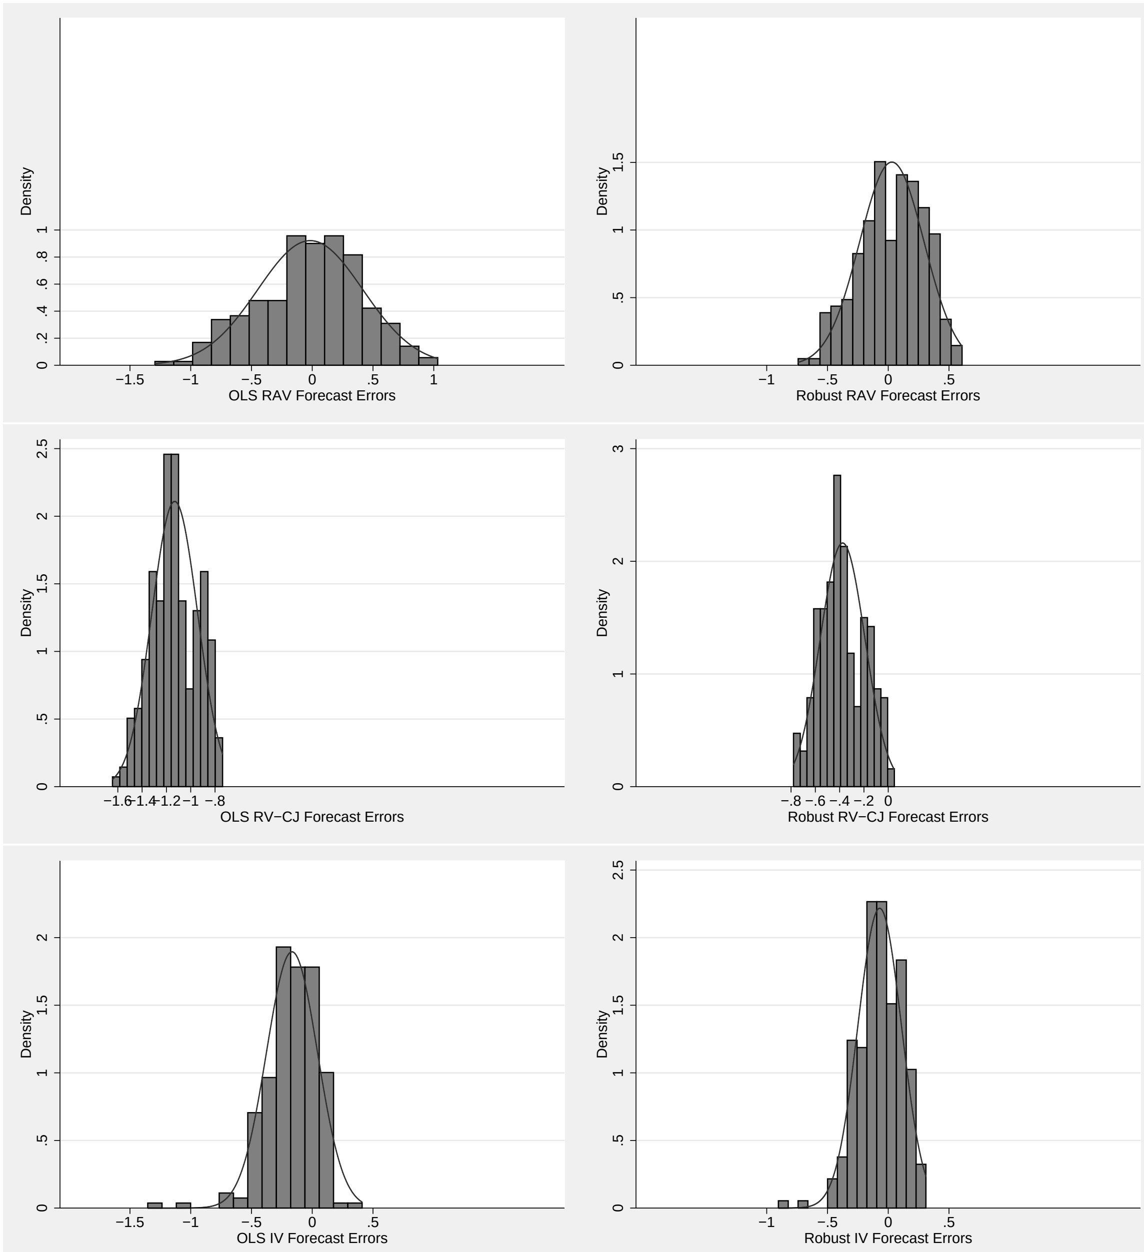

Figure 8: Plot of Actual Variance Against Different Forecasts for SPY

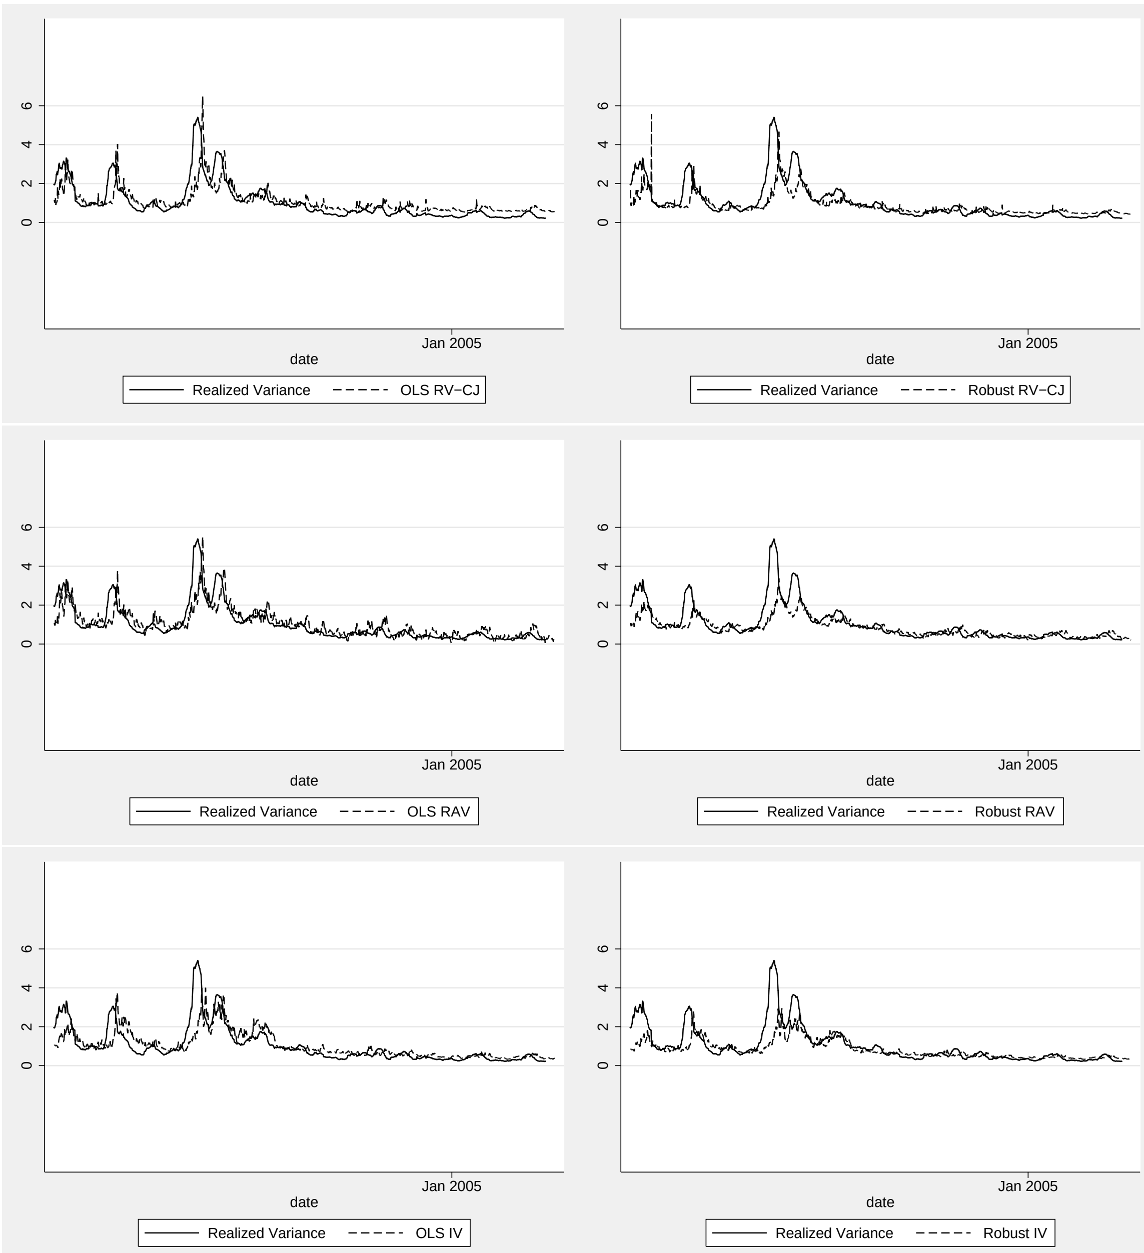

Figure 10: Histograms of Forecast Errors for 22 Day Forecasts of the Variance of the SPY

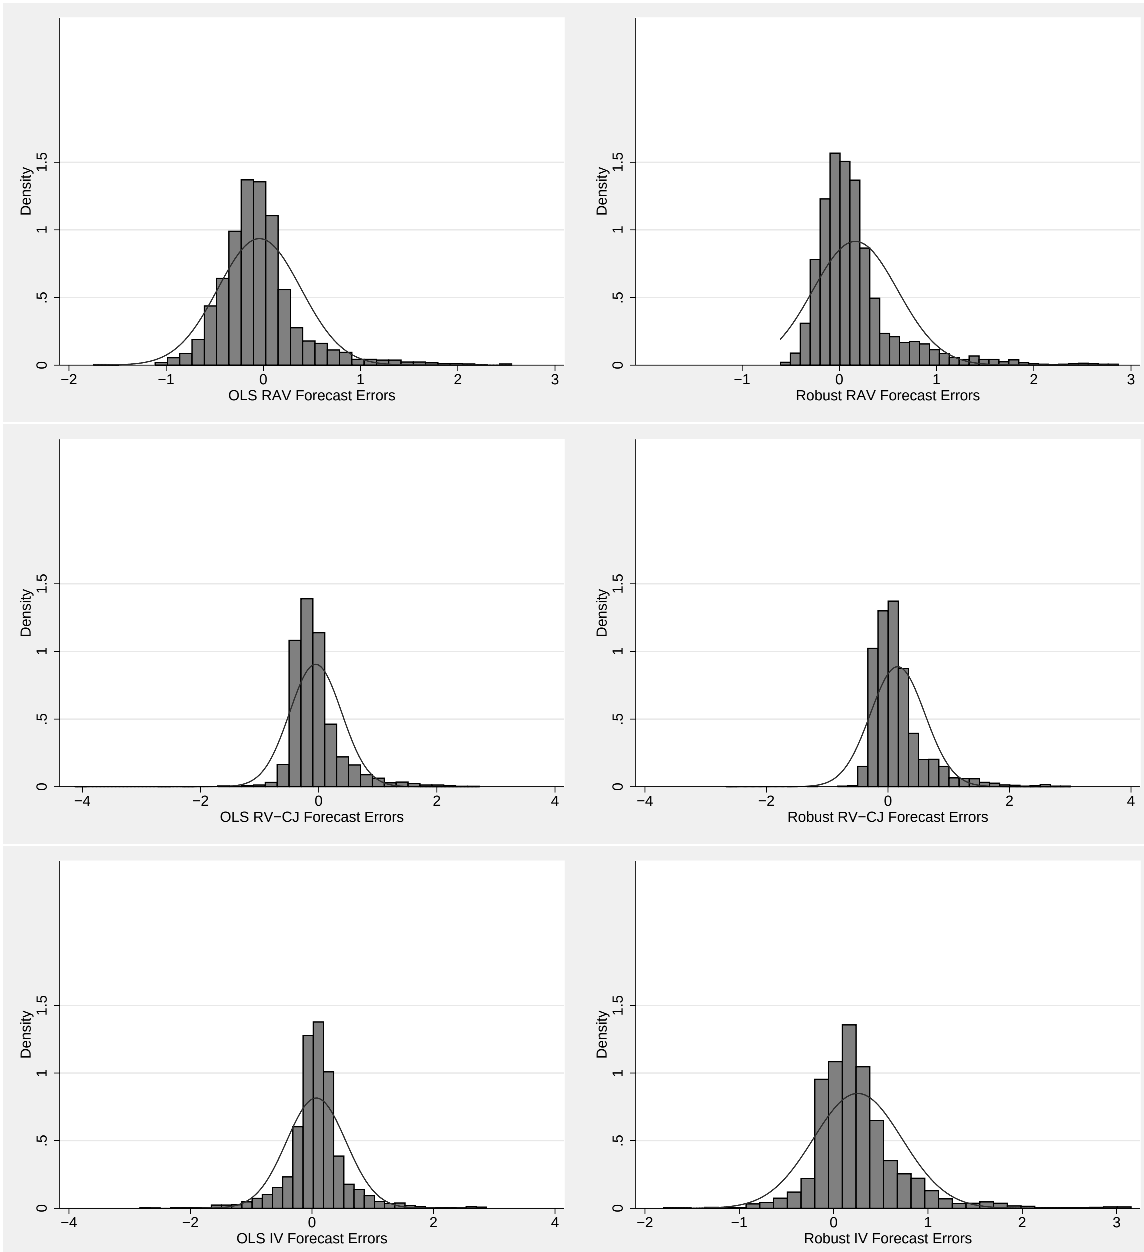

13 Technical Appendix

Additional Table 1: Summary Statistics											
	Variable	Mean	Std. Dev.	Min	Max		Variable	Mean	Std. Dev.	Min	Max
BMY	RV	3.17	4.11	0.27	45.28	MDT	RV	2.24	2.7	0.22	34.9
	Imp. Vol.	4.04	3.25	0.45	21.27		Imp. Vol.	2.89	1.86	0.59	12.29
	RAV	1.42	0.73	0.44	6.32		RAV	1.23	0.57	0.3	5.4
	Jump	0.04	0.27	0	6.7		Jump	0.05	0.56	0	17.33
C	RV	2.98	5.26	0.17	97.38	MOT	RV	6.99	8.89	0.34	175.37
	Imp. Vol.	3.6	3.4	0.1	34.81		Imp. Vol.	9.9	7.18	1.17	46.31
	RAV	1.39	0.84	0.3	9.29		RAV	2.17	1.04	0.41	10.05
	Jump	0.02	0.21	0	5.9		Jump	0.09	0.8	0	19.92
GE	RV	2.65	3.69	0.11	53.55	NOK	RV	4.17	4.98	0.2	43.54
	Imp. Vol.	3.54	2.92	0.4	17.24		Imp. Vol.	9.21	6.86	1.38	38.54
	RAV	1.34	0.75	0.23	6.21		RAV	1.69	0.91	0.32	6.41
	Jump	0.01	0.17	0	4.29		Jump	0.04	0.65	0	21.5
GS	RV	2.54	2.88	0.1	39.5	TXN	RV	7.8	8.35	0.51	62.92
	Imp. Vol.	4.3	3.35	0.8	20.99		Imp. Vol.	10.14	7.1	1.39	41.08
	RAV	1.36	0.61	0.23	4.58		RAV	2.33	1.14	0.51	8.07
	Jump	0.03	0.7	0	24.09		Jump	0.12	0.9	0	17.12
HD	RV	3.26	4.01	0.21	50.42	SPY	RV	1.03	1.3	0.05	16.25
	Imp. Vol.	4.43	3.42	1	20.56		Imp. Vol.	1.89	1.39	0.42	8.06
	RAV	1.49	0.74	0.26	6.06		RAV	0.85	0.43	0.16	3.54
	Jump	0.03	0.24	0	4.83		Jump	0.01	0.23	0	7.94
KO	RV	1.69	1.9	0.15	25.48	SPX	RV	1.03	1.3	0.05	16.25
	Imp. Vol.	2.13	1.56	0.29	11.41		Imp. Vol.	1.58	1.23	0.28	7.3
	RAV	1.11	0.49	0.25	4.98		RAV	0.85	0.43	0.16	3.54
	Jump	0.02	0.18	0	4.18						

Additional Table 2: Adjusted R^2 Values for OLS Regressions 1 Day Ahead

	BMY	C	GE	GS	HD	KO	MDT	MOT	NOK	TXN	SPY	SPX	AVG
HAR-RV	**0.49**	0.52	0.45	0.58	0.59	0.56	0.48	0.36	0.62	0.63	0.46	0.46	0.53
HAR-RV-CJ	0.48	0.52	0.45	0.62	0.60	0.56	**0.49**	0.36	**0.62**	0.63	0.49	0.49	0.53
HAR-RAV	0.48	0.53	0.48	0.59	0.61	0.56	0.49	**0.39**	0.61	0.63	0.50	0.50	0.54
HAR-RAV^2	0.50	**0.56**	0.48	**0.63**	**0.62**	**0.57**	0.48	0.38	0.62	**0.64**	**0.51**	**0.51**	**0.55**
Implied Volatility	0.45	0.52	**0.49**	0.56	0.56	0.50	0.44	0.35	0.55	0.62	0.42	0.43	0.50
HAR-RV-CJ + IV	0.50	0.54	0.50	0.63	0.61	0.58	0.51	0.39	0.63	0.66	0.49	0.50	0.56
HAR-RAV + IV	0.50	0.57	0.52	0.62	0.63	0.59	0.52	0.41	0.62	0.67	0.51	0.51	0.56
% From RV to CJ	0.00	0.00	0.00	0.03	0.01	0.00	0.01	0.00	0.00	0.00	0.02	0.02	0.01
% From RAV to CJ	-0.01	-0.03	-0.03	-0.01	-0.02	-0.01	0.01	-0.02	0.01	-0.01	-0.03	-0.03	-0.01
% From CJ to Combo	0.02	0.02	0.05	0.01	0.01	0.02	0.02	0.03	0.01	0.03	0.01	0.01	0.02
% From IV to Combo	0.06	0.03	0.02	0.06	0.05	0.08	0.07	0.04	0.08	0.04	0.07	0.07	0.05
% From RAV to Combo2	0.02	0.04	0.04	0.03	0.02	0.03	0.04	0.02	0.01	0.04	0.01	0.01	0.03
% From IV to Combo2	0.05	0.05	0.03	0.06	0.07	0.08	0.08	0.05	0.07	0.05	0.09	0.09	0.06

AVG is the average value over all 10 individual stocks

% Signifies Percentage Change in Adjusted R^2 from one model to another

Combo is a model where the independent variables from the HAR-RV-CJ Model are combined with Implied Volatility

Combo2 is a model where the independent variables from the HAR-RAV Model are combined with Implied Volatility

Bold values signify the highest adjusted R^2 in the non-combo models

Additional Table 3: Adjusted R^2 Values for OLS Regressions 5 Day Ahead

	BMY	C	GE	GS	HD	KO	MDT	MOT	NOK	TXN	SPY	SPX	AVG
HAR-RV	0.67	0.45	0.53	0.66	0.65	0.59	0.66	0.53	0.71	0.70	0.58	0.58	0.62
HAR-RV-CJ	0.67	0.45	0.54	0.69	0.65	0.59	0.67	0.53	0.72	0.71	0.60	0.60	0.62
HAR-RAV	0.67	0.52	0.60	0.68	0.70	0.62	0.67	0.58	0.72	0.72	0.65	0.65	0.65
HAR-RAV^2	0.68	0.47	0.56	0.68	0.66	0.60	0.66	0.56	0.71	0.71	0.61	0.61	0.63
Implied Volatility	0.67	0.53	0.60	0.68	0.66	0.55	0.64	0.52	0.66	0.71	0.54	0.53	0.62
HAR-RV-CJ + IV	0.73	0.54	0.62	0.73	0.71	0.64	0.71	0.58	0.74	0.75	0.63	0.63	0.68
HAR-RAV + IV	0.72	0.55	0.64	0.72	0.72	0.64	0.73	0.61	0.74	0.75	0.65	0.65	0.68
% From RV to CJ	0.00	0.00	0.01	0.03	0.01	0.00	0.01	0.00	0.01	0.00	0.03	0.03	0.01
% From RAV to CJ	-0.01	-0.02	-0.02	0.01	-0.01	0.00	0.01	-0.03	0.01	0.00	-0.01	-0.01	-0.01
% From CJ to Combo	0.06	0.09	0.08	0.04	0.06	0.05	0.05	0.05	0.03	0.05	0.03	0.03	0.05
% From IV to Combo	0.06	0.01	0.02	0.05	0.05	0.08	0.08	0.06	0.08	0.05	0.10	0.10	0.05
% From RAV to Combo2	0.05	0.03	0.04	0.04	0.03	0.03	0.06	0.03	0.02	0.04	0.01	0.01	0.04
% From IV to Combo2	0.05	0.02	0.04	0.03	0.07	0.09	0.09	0.08	0.08	0.05	0.12	0.12	0.06

AVG is the average value over all 10 individual stocks

% Signifies Percentage Change in Adjusted R^2 from one model to another

Combo is a model where the independent variables from the HAR-RV-CJ Model are combined with Implied Volatility

Combo2 is a model where the independent variables from the HAR-RAV Model are combined with Implied Volatility

Bold values signify the highest adjusted R^2 in the non-combo models

Additional Table 4: Adjusted R^2 Values for OLS Regressions 22 Days Ahead

	BMY	C	GE	GS	HD	KO	MDT	MOT	NOK	TXN	SPY	SPX	AVG
HAR-RV	0.60	0.38	0.44	0.64	0.45	0.43	0.53	0.42	0.61	0.59	0.45	0.45	0.51
HAR-RV-CJ	0.62	0.38	0.45	0.67	0.46	0.43	0.56	0.43	0.62	0.61	0.47	0.47	0.52
HAR-RAV	0.64	**0.51**	0.55	0.66	**0.53**	**0.49**	0.55	**0.52**	**0.68**	**0.64**	**0.56**	**0.56**	**0.58**
HAR-RAV^2	0.61	0.41	0.46	0.66	0.46	0.43	0.50	0.46	0.62	0.59	0.48	0.48	0.52
Implied Volatility	**0.71**	0.50	**0.55**	**0.68**	0.50	0.46	**0.58**	0.46	0.60	0.58	0.47	0.46	0.56
HAR-RV-CJ + IV	0.73	0.50	0.56	0.72	0.52	0.50	0.62	0.49	0.66	0.63	0.52	0.51	0.59
HAR-RAV + IV	0.73	0.53	0.59	0.70	0.55	0.51	0.63	0.54	0.69	0.65	0.57	0.57	0.61
% From RV to CJ	0.02	0.00	0.01	0.03	0.02	0.00	0.03	0.01	0.01	0.01	0.02	0.02	0.01
% From RAV to CJ	0.01	-0.03	-0.01	0.01	0.01	0.00	0.06	-0.03	0.00	0.02	-0.01	-0.01	0.00
% From CJ to Combo	0.11	0.12	0.10	0.05	0.06	0.07	0.06	0.06	0.03	0.03	0.05	0.04	0.07
% From IV to Combo	0.02	0.00	0.00	0.05	0.02	0.04	0.04	0.03	0.06	0.05	0.05	0.05	0.03
% From RAV to Combo2	0.10	0.02	0.03	0.04	0.02	0.02	0.08	0.02	0.01	0.01	0.01	0.01	0.03
% From IV to Combo2	0.02	0.03	0.03	0.02	0.04	0.05	0.05	0.08	0.09	0.07	0.10	0.11	0.05

AVG is the average value over all 10 individual stocks

% Signifies Percentage Change in Adjusted R^2 from one model to another

Combo is a model where the independent variables from the HAR-RV-CJ Model are combined with Implied Volatility

Combo2 is a model where the independent variables from the HAR-RAV Model are combined with Implied Volatility

Bold values signify the highest adjusted R^2 in the non-combo models