

Volatility Forecasting in the 90-Day Australian Bank Bill Futures Market

Nathan K. Kelly^a, J. Scott Chaput^b

^a Ernst & Young
Auckland, New Zealand

^b Lecturer
Department of Finance and Quantitative Analysis
University of Otago, Dunedin, New Zealand

Abstract

This study employs a comprehensive data set from the 90-Day Australian Bank Bills Futures market to test the predictive power of the theoretically superior implied volatility against historical volatility. Overall, the volatility forecasting results for are in line with previous research into futures markets. Implied volatility is a bias forecaster, historical volatility contains no extra information beyond that contained in implied volatility and the market is relatively efficient. Due to the use of both overlapping and non-overlapping data sets and the unique attributes to the market, futures style margining on the option contracts and geographic location, the results are significant.

Volatility Forecasting in the 90-Day Australian Bank Bill Futures Market

Volatility is the extent at which the return on an underlying asset fluctuates over a given period of time. It is most commonly calculated as the annualized standard deviation of returns and represents the risk associated with that asset. Historically, financial price series have shown great variation in volatility over time. Furthermore, there is significant evidence of volatility ‘clustering,’ periods of high volatility tend to occur together and likewise for periods of low volatility. As volatility represents risk, the clustering is important to market participants. However, on a given day, past volatility itself is not necessarily important, what is important, is future volatility. This is because volatility is a key component of many financial decisions, including asset allocation, risk management and portfolio selection.

This paper tests the forecasting ability of two distinctively different measures of volatility: implied and historical. The market that is tested is the 90-Day Australian Bank Bill Futures market. The general procedure for tests of volatility forecasting is to compare the implied volatility (IV) and historical volatility (HV) as forecasts, with the subsequent or realized volatility (RV) that relates to the forecast period. This study is consistent with this method.

The rest of the paper is organised as follows. A brief overview of volatility forecasting follows immediately. After that is a comprehensive review of previous research and the justification for testing volatility forecasting in the 90-Day Australian Bank Bill Futures market. The data is then described and the methodology of the research set out, including a discussion of the various models for calculating and

obtaining volatility forecasts. The statistical test procedures are then discussed in detail and, lastly, the empirical results are presented and conclusions drawn.

Volatility Forecasting

Market participants and academics have long been aware of the significance of predicting volatility. Early development of volatility forecasting theory relied on using the historical standard deviation of returns as a measure of an asset's risk. However, it was soon realised that an asset's volatility was not constant over its lifetime. Thus, it became important to provide accurate measures of volatility that captured this changing effect. This led rise to numerous theories for calculating. These measures of volatility are moving averages (MA), exponentially weighted moving averages (EWMA), autoregressive conditional heteroskedastic models (ARCH), and generalised autoregressive heteroskedastic models (GARCH). It is important to realise all of these models rely on past data in an attempt to predict future volatility. There is, however, one method of volatility forecasting that is forward looking, implied volatility.

The famed Black-Scholes-Merton (BSM) option pricing model has only one parameter that is not directly observable in the market, the underlying asset's volatility. When there is an active options market on an asset, the prices of these options are based on the market participants' expectation of future volatility. Therefore, by observing an option's price and using the BSM option pricing model you can solve for, or 'back out', the underlying asset's IV. This is the market's unbiased expectation of the asset's volatility for the time period until the option matures, given the assumption that the option pricing model is correct.

IV should be a superior forecaster because it is effectively traded through options. If a trader feels volatility will increase in the future, she can buy options. If she feels volatility will decrease, she can sell options. By maintaining a delta-neutral hedge, she can capture the difference between realized volatility and IV. IV is a forward looking measure of expected future volatility. Therefore, IVs change as traders react to new information. IV has a degree of forward looking that HV does not. IV can be thought of as being derived from a superior information set. If market participants are efficiently processing new information available to them, they will form expectations of volatility that are, on average, correct, and at least as good as historical measures. This is in line with the unbiased expectations hypothesis and, subsequently, the test that IV is a superior forecaster compared to HV has the joint hypothesis of market efficiency.

Prior Research

Early research into volatility forecasting indicated that IV had superior predictive power compared to HV. Latane and Rendleman (1976) used stock options and calculated weighted implied standard deviations (WISDs), an average of IVs from multiple options, differing only in strike price. They concluded that “the WISD is generally a better predictor of future volatility than standard deviation predictors based on historical data.” However, Latane and Rendleman (1976) ignored many important considerations such as dividend payments, transaction costs, taxes and timing differences between stock and option markets that could have affected their results. They also found that volatility is not constant over time and that individual stock volatilities tend to move together. A key assumption of the BSM option pricing

model is that the volatility is constant over the life of the option or changes in a known way and Latane and Rendleman (1976) found evidence against this.

Chiras and Manaster (1978) used Merton's generalised form of the Black-Scholes model that accounts for dividends and still found IV to be better predictors of subsequent stock return variance than those obtained from historical data. They also devised a trading strategy to exploit this and found it provided abnormally high returns. Using stock options, Schmalensee and Trippi (1978) and Beckers (1981) also find support for IV and the investors' ability to make accurate forecasts.

More recent work found contrasting results to the initial findings. Canina and Figlewski (1993)¹ (CF) argue that Beckers (1981) was the only researcher to identify the importance of timing. The other studies relied on monthly data and did not match RVs with the time to maturity remaining on the option. They argue that many of the papers even analyzed forecasts of volatility over periods before the forecast was made. This would clearly lead to incorrect results.

CF tested for forecasting ability using S&P 500 index options, but their test differed in a significant way. All previous research has relied on the assumption that the BSM model is appropriate and accurate. However, CF recognised that this is not the case for American style options, the most common type of exchange-traded options. Therefore, they used a binomial model to capture the value of early exercise. They argue that this accounts for the possibility that the market may not price the options with exactly the same model as the researcher does. CF used regression tests and found that IV is a poor forecaster of subsequent RV and that it did not even incorporate the information present in HV. Further supporting CF, Lamoureux and

¹ Canina and Figlewski (1993) is perhaps the most well known paper that tests the forecasting ability of various volatility estimates. This is because of the regression methodology they developed. A very similar methodology is adopted for this paper and it is discussed in more detail later in the paper.

Lastrapes (1993) used individual equity options and rejected the hypothesis that available information cannot be used to improve the markets forecast, IV.

However, while CF's paper is highly regarded it does not escape criticism. Corrado and Miller (2005) (CM) perform the tests similar to CF with a larger data set that includes CF's data. They find that the forecasting ability of IV has improved since 1995 and that it is now superior to HV.

CM split their data into two basic periods, pre-1995 and post-1995. They find that adding any independent variables beyond IV does not appear to add any more explanatory power the model. This means that that all the information in historical measures is consumed by IV measures. Therefore, they conclude that IV is a better forecaster than HV.

CM also suggest the results obtained by studies such as CF are contaminated by methodological issues and econometric problems. The major concern of CM is there are significant biases present when the forecasting ability of IV is tested. The main ones are the errors-in-variables problem and model misspecification bias. Errors-in-variables mean that the observed values of the independent variables include errors. Their paper tests this by running instrumental variable regressions and comparing them to the OLS regression. They find that the errors-in-variables problem drops out after 1995. The model misspecification bias occurs when you must decide how to calculate HV. A simple standard deviation of returns is used in most studies, while there are arguments that various models which weight past volatilities differently should in fact be used. Their paper tests this by using various models and finds that there is no benefit gained by using the more complex models compared to more simple methods. This suggests that using the much vaunted ARCH, GARCH

and EWMA models are no better than using the simpler standard deviation of past returns, when testing volatility forecasting.

Forecasting Volatility in Futures Markets

So far, all previous research that has been discussed has focused on forecasting volatility in equity markets. An important implication of the results from these tests is that they are specific to each data set. One cannot conclusively infer from them the best method for forecasting volatility in any market, as there may be specific market characteristics that cause results to differ.

Szakmary et al. (2003) provide a comprehensive test of 35 futures markets from the U.K. and U.S. The types of markets tested include equities, currency, interest rate and commodity futures markets. In 34 out of the 35 markets tested, they find IV to be a superior forecaster compared to HV, with sugar the sole exception.

Szakmary et al. (2003) argue that the reason the results they find are so conclusively in favour of market efficiency is due to the nature of futures markets. The key characteristic that futures markets have is that both the futures contract and the option contract are traded on the same exchange, whereas this is not the case for equity markets. This minimises transaction costs and reduces the barriers to arbitrage. Also, because they are traded on the same exchange, the two markets close at the same time. Equity markets close fifteen minutes prior to options markets, meaning that the closing price of the asset does not correspond to the closing price of the option. Therefore, the IV calculated may not represent the market's unbiased expectation of future volatility. This non-synchronicity could introduce measurement error to the independent variables. Corrado & Miller (2005) identified this as the errors-in-variables problem and suggest that it could lead to biased results. However,

Szakmary et al. (2003) argue this problem is overcome when the underlying market being tested is a futures market and the use of instrumental variables is not needed.

Data and Methodology

Description of data

We look at the 90-Day Australian Bank Bill Futures and their options. Our data are from the Sydney Futures Exchange and cover the period of September 1992 to December 2000. The 90-Day Australian Bank Bill Futures series contains settlement prices for the futures with expiration in March, June, September and December, with term to maturities up to five years. The Options price series contains put and call prices for varying strike prices. The term to maturities for the options up to two years and the options expire seven days before the underlying futures contract expire. The options are American and have futures style margining. The investor pays the full premium to purchase the option and gains and losses are marked-to-market daily.

We create two distinctive data sets. The first is an overlapping data series, where there is an IV, HV, and RV value for every day possible. It is overlapping in the sense that on consecutive days the HV and RV values will be made up of largely the same observation windows of futures returns. The final overlapping data series contains 10608 observations. The second data set used for this research is three non-overlapping series. This involves using only one set of observations from each futures contract, determined by a specific time to maturity. The times to maturity used are 21, 41, and 62-trading days and this limited the number of observations to 33 or 32 for each series. These time to maturities are chosen to reflect 30, 60, and 90-calendar day

windows respectively and the method is consistent with Jorion (1995), Christensen and Prabhala (1998), and Szakmary et al. (2003) to allow comparisons to be made.

The reason for using the two data series is to add to the significance and comparability of the results. The overlapping series has a very large number of observations that make the regression tests very powerful. However, due to the overlapping nature of the observation windows, there is an inherent level of serial correlation in the HV and RV series. Therefore, following Christensen and Prabhala (1998) and Corrado and Miller (2005), the Newey-West (1987) correction method is used to adjust the standard errors of the regression coefficients to accurately reflect the serial correlation of varying lengths in the residuals². The non-overlapping series, on the other hand, has the benefit that each observation of RV and HV contains no overlapping futures return data. This means that there is no underlying serial correlation in the data.

Also, by testing both sets of data this study inadvertently tests whether or not there is any difference between the two techniques and allows the results to be comparable with as many previous studies as possible. A formal presentation of the models used to calculate IV, HV and RV follows. There is also a complete example of the methodology provided in the appendix for further clarification.

Volatility Estimation

Implied Volatility

The model that is used to price options with futures style margining is an extension of the Black (1976) model for options on futures contracts and was

² See Newey-West (1987) for details of the method.

developed by Asay (1982). It can be thought of as the forward price of Black's model.

Formally, the Asay model for the call premium is:

$$C = FN(d_1) - XN(d_2)$$

Where:

$$d_1 = \frac{\ln(F/X) + \frac{1}{2}\sigma^2 t}{\sigma\sqrt{t}}$$

$$d_2 = d_1 - \sigma\sqrt{t}$$

And F = futures price;
X = exercise price;
C = call price;
t = time to maturity, and;
 σ = instantaneous volatility.

The put premium is given by:

$$P = XN(-d_2) - FN(-d_1)$$

Lieu (1990) showed that when options have futures style margining there is no value to early exercise. Therefore, this eliminates a common source of bias in IV's that may have influenced previous results. For these reasons, it would be expected that the Asay (1982) model is accurate in the Australian short-term interest rate market for near-the-money options.

As mentioned, if the model is not pricing the options accurately then the IV will be a biased forecast. Therefore, as the Asay model is most likely to be accurate for the nearest-to-the-money options, only these are used to calculate the implied volatility. This consists of using the nearest-to-the-money call and the nearest-to-the-money put. For the S&P 100, Harvey and Whaley (1991) showed that the put and call IVs are negatively correlated, so a more efficient estimator is obtained by combining

them³. In this case, because of the relatively small financing and transaction costs for futures options relative to cash instrument options, the IV's of the puts and calls are almost identical; however, they are still combined to ensure the accuracy of the forecast. Secondly, all options contracts with a term to maturity of less than 10 trading days were removed because their IVs have been shown to be unstable⁴.

A Newton-Raphson method was employed to back out the IV from each option price. This is because the option premium is a non-linear function of the volatility, so IV could not be solved for without a numerical estimation method. This gives us one IV forecast for each option price. Following CF, the first step to test market efficiency is to recognize that IV is regarded as a forecast of the market's RV. To formally state this:

$$IV = E_{MKT} [RV]$$

Where, E_{MKT} is the market's expectation of the RV.

Historical Volatility

The HV model that is used as the benchmark in this paper is the sample standard deviation of daily futures returns. Previous research indicates that the more complex models for measuring HV such as ARCH, GARCH and EWMA have added no significant explanatory power to forecasting tests compared to more simple measures⁵. Therefore, for this research, only the historical annualized standard deviation of returns is used as a measure of HV.

To determine the matching HV forecast for each RV, a historical window that matched the time to maturity of the option was used. This is Figlewski's optimized unconditional conditional heteroskedastic (OUCH)⁶ estimator. Figlewski found that

³ Corrado and Miller (2005)

⁴ Szakmary et al. (2003)

⁵ Corrado and Miller (2005) and Szakmary et al. (2003)

⁶ Figlewski (1997)

the accuracy of the forecast was highest when the window, from which the historical measure of volatility was calculated, equalled the time to maturity of the option. Formally, the HV estimate for a particular day is calculated as the annualized standard deviation of returns on the futures contract from the previous TM days, where TM is the term to maturity of the option contract used to calculate the IV observation on that same day. This only applies to the full overlapping data set, HV. In the non-overlapping series, the HV observation represents the RV observation from the previous futures contract. This is done to ensure consistency with previous research. This means that for observations in the non-overlapping data series, HV21, HV41, and HV62, the values are simply lags of the relevant RV series, RV21, RV41, and RV62.

Realized Volatility

The RV for each day is calculated as the annualized standard deviation of returns of the futures contract from the observation date until the option contract that the IV estimate is obtained from expires. It is essentially the forward looking version of each HV observation. This means that the observation dates relating to the futures contracts whose options expired after December 2000 had to be removed as the RV would be unable to be calculated. Formally, following Canina and Figlewski (1993) and Szakmary et al. (2003), the RV is calculated as the annualized standard deviation of the continuously compounded futures returns:

$$RV = \left(\frac{250}{T_M - 1} \sum_{t=1}^{T_M} (R_t - \bar{R})^2 \right)^{1/2}$$

Where,

TM is the time to maturity in days of the option contract,

$R_t = \ln (P_t/P_{t-1})$, the return on the futures contract, and;

\bar{R} is the sample mean of R_t .

Figure 1 provides a timeline of the volatility estimation periods.

Descriptive Statistics

The descriptive statistics of each data series are shown in Table 1. The mean and standard deviation of all the series is relatively similar. Also, the range of each series is similar with the RV having the largest. The values of note are the relatively high values of skewness and kurtosis, especially for the HV21 series. There is a possibility that this will influence the regression results for this data set. Overall, however, the data is relatively similar to previous research and appears adequate to warrant further analysis.

The subsequent methodology used to test forecast power is that set out by CF and is consistent with Szakmary et al (2003) in their tests of futures markets. It comprises the standard tests for the suitability of a times series for regression analysis and three regressions themselves. These statistical procedures are discussed in detail shortly, but first, the research hypotheses are presented.

Research Hypotheses

IV is regarded as an unbiased expectation of the RV under the assumption that the market is informationally efficient (weak-form market efficiency) and the option pricing model is specified correctly. If the Asay pricing model for options with futures style margining is accurate for the short-term near-the-money options used in this study, then rejection of the IV as superior predictor compared to HV would imply that the market does not efficiently price options. This would mean that the Australian

short-term interest rate market is not efficient and that arbitrage opportunities could exist.

Consistent with the existing literature the following three hypotheses are examined⁷:

H1: IV is an unbiased estimate of the future RV.

H2: IV has more explanatory power than the HV in forecasting RV.

H3: IV efficiently incorporates all information regarding future volatility;
HV contains no information beyond what is already included in IV.

These hypotheses are in order of power. Failure to reject H1 will mean failure to reject H2 and H3. However, failure to reject H3 will not mean failure to reject H1 or H2. The next section describes the regressions we use.

Regression Tests

The procedure to test the forecasting ability of IV versus HV in CF and Szakmary et al. (2003) is employed by estimating the following regressions:

Unbiased Predictor Regressions

$$RV_t = \alpha_0 + \beta_1 IV_t + \varepsilon_t \quad (1)$$

$$RV_t = \alpha_1 + \beta_2 HV_t + \varepsilon_t \quad (2)$$

Encompassing Regression

$$RV_t = \alpha_0 + \beta_1 IV_t + \beta_2 HV_t + \varepsilon_t \quad (3)$$

Consistent with hypothesis H1, IV is an unbiased predictor of the RV, it is expected that $\alpha_0 = 0$ and $\beta_1 = 1$ in regression (1).

$$H1: \quad \alpha_0 = 0, \beta_1 = 1$$

⁷ These hypotheses are taken from Szakmary et al. (2003).

H1A: IV is a biased predictor of RV

Consistent with hypothesis H2, IV includes more information (i.e., current market information) than HV, then IV should have greater explanatory power than HV, and it would be expected that regression (1) has a higher adjusted R^2 than regression (2).

H2: $\text{adj}R^2(1) > \text{adj}R^2(2)$

H2A: HV contains more information about future volatility than IV

Finally, consistent with hypothesis H3, when IV and HV appear in the same regression, as in (3), it would be expected that $\beta_2 = 0$ since HV should have no explanatory power beyond that already contained in IV. The adjusted R^2 ($\text{adj} R^2$) should also be the same for regressions (1) and (3) if no extra explanatory power is gained from HV. Adjusted R^2 is used to measure extra explanatory power as it is well known that adding any independent variable increases R^2 .

H3: $\beta_2 = 0, \text{adj}R^2(3) = \text{adj}R^2(1)$

H3A: HV contains information about future volatility not found in IV

The results of the statistical procedures and hypothesis test are formally discussed in the following section.

Empirical Results

Forecasting Regression Results

The results of the tests for the relative predictive power of IV and HV appear in Table 2. The first three regressions correspond to the overlapping data set of 10608 observations. The next three correspond to the 21 trading day non-overlapping series, the following three to the 41 trading day non-overlapping series and the last three regressions relate to the 62 day non-overlapping series. For each regression, the

coefficients and adjusted R²s are shown, with relevant t-stats reported in brackets below the estimate. The results of the relevant hypothesis tests are shown in Table 3.

Unbiased Predictor Regressions

Regressions (1), (4), (7) and (10) represent the tests of whether or not IV is an unbiased predictor of RV. The coefficients for IV range from 0.438 to 0.583, and all of the intercepts are significantly greater than zero. If IV is an unbiased predictor of RV, as H1 states, then the intercept should not be significantly different from zero and the coefficient for IV should not be significantly different than one. This is clearly not the case in any of the series, with the hypothesis rejected in all of the regressions. IV is not an unbiased predictor of RV in any of the series tested. This result is consistent with Jorion (1995), Szakmary et al. (2003) and other previous studies, where the coefficient on IV is found to be between 0.300 and 0.650. The slope coefficients for IV are significantly greater than zero in all of the regressions except for (7), the 21 trading day non-overlapping series. The adjusted R²'s from the regressions range from 0.0828 for the 21-trading day series to 0.2341 for the 62-trading day series, with the other two close to the 0.220. This indicates that IV is useful in predicting RV. However, the results are comparable with previous research into interest rate futures markets, where R² values ranges from 0.106 to 0.5008. Therefore, on a relative basis, the results of IV's ability to predict RV in the 90-Day Australian Bank Bills Futures market are in line with world markets.

Regressions (2), (5), (8) and (11) represent the tests of the predictive power of HV. The slope coefficients for HV range from 0.23127 to 0.38145, with only one significantly greater than zero. This was for regression (2), the full overlapping series. HV does not have any significant forecasting power using non-overlapping. Further

⁸ Szakmary et al. (2003)

illustrating this point is that the intercept term is positive and highly significant in all four regressions. The adjusted R^2 's for HV range from 0.0199 to 0.0827. This is very low relative to previous research and makes it quite clear that HV is not useful in predicting RV in this market. Hypothesis two, IV has more explanatory power than HV, required that the adjusted R^2 's from regressions (1), (4), (7), and (10) be higher than those from regressions (2), (5), (8), and (11) respectively. This is true in every instance, with the difference ranging from 0.0620 to 0.1969. These results show IV has more explanatory power than HV in forecasting RV in the 90-Day Australian Bank Bills Futures market, can be accepted.

Encompassing Regressions

Regressions (3), (6), (9) and (12) represent the encompassing regression tests. They are used to directly compare the predictive power of IV versus HV. If IV contains all of the information that HV does plus any more relevant information then it would be expected that the coefficient on HV to fall to zero in the encompassing regressions and the adjusted R^2 's to be no greater than those for the regression with IV as the only independent variable. The coefficients on HV range from 0.03159 to 0.08024, with the only significant result in the overlapping series. The adjusted R^2 's from the encompassing regressions are comparable to the regressions where IV is the only independent variable. This indicates that adding HV to the model provides no extra predictive power above that gained from IV. Therefore, hypothesis three, the IV efficiently incorporates all information regarding future volatility and that HV contains no information beyond that contained in IV, can be accepted. The coefficients for IV are significant in three out of the four regressions, with a range from 0.43383 to 0.56832, similar to results for IV alone. Again, the only series where

the coefficient of IV is not significant is the 21-day non-overlapping series. All of the coefficients for IV are significantly closer to one than those for HV and the t-stats ranges from three to over twenty times larger for IV.

Ideally, if IV was an unbiased forecaster of RV and contained all relevant forecasting information the intercept term and coefficient for HV would not be significantly different from zero and the coefficient for IV would not be significantly different from one. This hypothesis is rejected in three of the regressions with the exception being the 21-trading day non-overlapping series. While it would be easy to assume that this is a standout result for good reasons; it is more likely the result of the unique statistical properties of this series relative to the rest. This suspicion is confirmed by the low R2 value of 0.0753, the lowest of all the regressions where IV is an independent variable. Overall, it is clear that IV is a biased estimator of RV.

Summary and Conclusions

In Summary, this study employs a comprehensive data set from the 90-Day Australian Bank Bills Futures market to test the predictive power of the theoretically superior implied volatility against historical volatility. Overall, the volatility forecasting results for are in line with previous research into futures markets. Implied volatility is a bias forecaster, historical volatility contains no extra information beyond that contained in implied volatility and the market is relatively efficient. Due to the use of both overlapping and non-overlapping data sets and the unique attributes to the market, futures style margining on the option contracts and geographic location, the results are significant.

References

- Beckers, S., 1981. Standard deviations implied in option prices as predictors of futures stock prices variability. *Advances in Futures and Options Research* 5 (1), 297-308.
- Black, F., 1976. The pricing of commodity contracts. *Journal of Financial Economics* 3 (1/2), 167-179.
- Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* 81 (3), 673-654.
- Canina, L., and Figlewski, S., 1993. The informational content of implied volatility. *Review of Financial Studies* 6 (3), 659-681.
- Corrado, C. and Miller, T., 2005. The forecast quality of CBOE implied volatility indexes. *The Journal of Futures Markets* 25 (4), 339-373.
- Day, T.E. and Lewis, C.M., 1993. Forecasting futures market volatility. *Journal of Derivatives*, 33-50.
- Figlewski, S., 1997. Forecasting Volatility. *Financial Markets, Institutions and Instruments* 6, 1-88.
- Hansen, L.P., 1982. Large sample properties of generalized method of moments estimators. *Econometrica* 50 (4), 1029-1054.
- Jackwerth, J.C. and Rubenstein, M., 1996. Recovering probability distributions from option prices. *Journal of Finance* 51 (5), 1611-1631.
- Jorion, P., 1995. Predicting volatility in the foreign exchange market. *Journal of Finance* 50 (2), 507-528.
- Kumar, R., Shastri, K., 1990. The predictive ability of stock prices implied in option premia. *Advances in Futures and Options Research* 4 (1), 165-176.
- Lamoureux, C.G., Lastrapes, W.D., 1993. Forecasting stock-return variance: Toward an understanding of stochastic implied volatilities. *Review of Financial Studies* 6 (2), 293-326.
- Latane, H.A., Rendleman Jr., R.J., 1976. Standard deviations of stock price ratios implied in option prices. *Journal of Finance* 31 (2), 369-381.
- Newey, D.B., West, K., 1987. A simple positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55, 703-708.
- Schmalensee, R., Trippi, R.R., 1978. Common stock volatility expectations implied by option premia. *Journal of Finance* 33 (1), 129-147.
- Sims, C.A., Stock, J.H., Watson, M.W., 1990. Inference in linear time series models with some unit roots. *Econometrica* 58 (1), 113-144.
- Szakmary, A., Ors, E., Kim, J.K., and Davidson III, W.N., 2003. The predictive power of implied volatility: Evidence from 35 futures markets. *Journal of Banking and Finance* 27, 2151-2175.

TABLE 1
Descriptive Statistics

Panel A: Overlapping Data						
	RV	IV	HV			
<i>N</i>	10608	10608	10608			
Mean	0.173683	0.182278	0.201012			
Std. Dev.	0.050284	0.040791	0.034940			
Min	0.009858	0.032612	0.041887			
Max	0.495051	0.351655	0.394674			
Skewness	-0.097878	-0.461940	0.134627			
Kurtosis	0.984226	0.400898	2.906848			
Panel B: Non-Overlapping Data						
	IV21	HV21	IV41	HV41	IV62	HV62
<i>N</i>		33		33		32
Mean		0.170756		0.190533		0.141527
Std. Dev.		0.048699		0.061199		0.050338
Min		0.067886		0.102927		0.062922
Max		0.275745		0.344936		0.278259
Skewness		2.079366		0.995257		1.172837
Kurtosis		6.290270		2.454742		1.932901

A positive value of Skewness indicates the distribution is skewed to the right and a negative value indicates a left skew.

The Kurtosis statistic represents excess kurtosis. A value above zero indicates a leptokurtic distribution (high peak and fat tails) and a negative value indicates a platykurtic distribution (smaller peak and thin tails).

The descriptive statistics for the non-overlapping series of RV are not given because they are the same as those for the non-overlapping series of HV. This is because the non-overlapping series of HV represents that lag of the non-overlapping series of RV.

RV_t represents the overlapping series of realized volatility, IV_t represents the overlapping series of implied volatility, and HV_t represents the overlapping series of historical volatility. RV_{21_t} , RV_{41_t} , and RV_{62_t} represent the non-overlapping series of realized volatility for 21, 41 and 62 trading days respectively. IV_{21_t} , IV_{41_t} , and IV_{62_t} represent the non-overlapping series of implied volatility for 21, 41 and 62 trading days respectively. HV_{21_t} , HV_{41_t} , and HV_{62_t} represent the non-overlapping series of historical volatility for 21, 41 and 62 trading days respectively.

TABLE 2
Forecasting Regression Results

	Model	Coefficient Estimates			$adjR^2$
		α_0	β_1	β_2	
(1)	$RV_t = \alpha_0 + \beta_1 IV_t + \varepsilon_t$	0.06749** (34.25)	0.58259** (55.23)		0.2233
(2)	$RV_t = \alpha_0 + \beta_2 HV_t + \varepsilon_t$	0.09701** (35.29)		0.38145** (28.31)	0.0702
(3)	$RV_t = \alpha_0 + \beta_1 IV_t + \beta_2 HV_t + \varepsilon_t$	0.06374** (24.37)	0.56832** (45.79)	0.03159* (2.18)	0.2236
(4)	$RV_{21_t} = \alpha_0 + \beta_1 IV_{21_t} + \varepsilon_t$	0.03401 (0.93)	0.55248 (1.97)		0.0828
(5)	$RV_{21_t} = \alpha_0 + \beta_2 HV_{21_t} + \varepsilon_t$	0.07902** (3.57)		0.23127 (1.29)	0.0208
(6)	$RV_{21_t} = \alpha_0 + \beta_1 IV_{21_t} + \beta_2 HV_{21_t} + \varepsilon_t$	0.02774 (0.74)	0.54951 (1.66)	0.08024 (0.41)	0.0753
(7)	$RV_{41_t} = \alpha_0 + \beta_1 IV_{41_t} + \varepsilon_t$	0.05567* (2.38)	0.43831** (3.10)		0.2168
(8)	$RV_{41_t} = \alpha_0 + \beta_2 HV_{41_t} + \varepsilon_t$	0.09389** (3.75)		0.23532 (1.27)	0.0199
(9)	$RV_{41_t} = \alpha_0 + \beta_1 IV_{41_t} + \beta_2 HV_{41_t} + \varepsilon_t$	0.04843 (1.71)	0.43383* (2.69)	0.06618 (0.37)	0.1935
(10)	$RV_{62_t} = \alpha_0 + \beta_1 IV_{62_t} + \varepsilon_t$	0.04313 (1.40)	0.58059** (3.24)		0.2341
(11)	$RV_{62_t} = \alpha_0 + \beta_2 HV_{62_t} + \varepsilon_t$	0.09068** (3.37)		0.3525 (1.93)	0.0827
(12)	$RV_{62_t} = \alpha_0 + \beta_1 IV_{62_t} + \beta_2 HV_{62_t} + \varepsilon_t$	0.03003 (0.89)	0.55136* (2.64)	0.13450 (0.74)	0.2393

RV_t represents the overlapping series of realized volatility, IV_t represents the overlapping series of implied volatility, and HV_t represents the overlapping series of historical volatility. RV_{21_t} , RV_{41_t} , and RV_{62_t} represent the non-overlapping series of realized volatility for 21, 41 and 62 trading days respectively. IV_{21_t} , IV_{41_t} , and IV_{62_t} represent the non-overlapping series of implied volatility for 21, 41 and 62 trading days respectively. HV_{21_t} , HV_{41_t} , and HV_{62_t} represent the non-overlapping series of historical volatility for 21, 41 and 62 trading days respectively.

Regressions (1), (4), (7), and (10) test the predictive power of IV. Regressions (2), (5), (8), and (11) test the predictive power of HV. Regressions (3), (6), (9), and (12) are encompassing regression tests to directly compare the predictive power of IV versus HV.

t -stats are reported in brackets below the relevant parameter estimate and follow the normal t distribution with $N-k$ degrees of freedom.

* and ** denote significance at the 5% and 1% level respectively.

TABLE 3
Hypothesis Test Results

	Model	Null Hypothesis		
		$\alpha_0 = 0,$ $\beta_i = 1$	$\beta_i = 1$	$\alpha_0 = 0,$ $\beta_1 = 1,$ $\beta_2 = 0$
(1)	$RV_t = \alpha_0 + \beta_1 IV_t + \varepsilon_t$	982.40**	1565.79**	
(2)	$RV_t = \alpha_0 + \beta_2 HV_t + \varepsilon_t$	2738.56**	2107.28**	
(3)	$RV_t = \alpha_0 + \beta_1 IV_t + \beta_2 HV_t + \varepsilon_t$			656.75**
(4)	$RV_{21_t} = \alpha_0 + \beta_1 IV_{21_t} + \varepsilon_t$	3.09*	2.55	
(5)	$RV_{21_t} = \alpha_0 + \beta_2 HV_{21_t} + \varepsilon_t$	9.17**	18.33**	
(6)	$RV_{21_t} = \alpha_0 + \beta_1 IV_{21_t} + \beta_2 HV_{21_t} + \varepsilon_t$			1.53
(7)	$RV_{41_t} = \alpha_0 + \beta_1 IV_{41_t} + \varepsilon_t$	16.55**	15.74**	
(8)	$RV_{41_t} = \alpha_0 + \beta_2 HV_{41_t} + \varepsilon_t$	8.53**	16.98**	
(9)	$RV_{41_t} = \alpha_0 + \beta_1 IV_{41_t} + \beta_2 HV_{41_t} + \varepsilon_t$			8.99**
(10)	$RV_{62_t} = \alpha_0 + \beta_1 IV_{62_t} + \varepsilon_t$	8.43**	5.47*	
(11)	$RV_{62_t} = \alpha_0 + \beta_2 HV_{62_t} + \varepsilon_t$	6.69**	13.33**	
(12)	$RV_{62_t} = \alpha_0 + \beta_1 IV_{62_t} + \beta_2 HV_{62_t} + \varepsilon_t$			4.63**

The test statistics reported are F -values. Critical values follow the usual F -distribution.

The null hypothesis: $\alpha_0 = 0, \beta_i = 1$, represents the test for an unbiased forecaster of RV. The null hypothesis: $\alpha_0 = 0, \beta_1 = 1, \beta_2 = 0$, tests whether IV contains all the predictive information of HV.

RV_t represents the overlapping series of realized volatility, IV_t represents the overlapping series of implied volatility, and HV_t represents the overlapping series of historical volatility. RV_{21_t} , RV_{41_t} , and RV_{62_t} represent the non-overlapping series of realized volatility for 21, 41 and 62 trading days respectively. IV_{21_t} , IV_{41_t} , and IV_{62_t} represent the non-overlapping series of implied volatility for 21, 41 and 62 trading days respectively. HV_{21_t} , HV_{41_t} , and HV_{62_t} represent the non-overlapping series of historical volatility for 21, 41 and 62 trading days respectively.

Regressions (1), (4), (7), and (10) test the predictive power of IV. Regressions (2), (5), (8), and (11) test the predictive power of HV. Regressions (3), (6), (9), and (12) are encompassing regression tests to directly compare the predictive power of IV versus HV.

* and ** denote significance at the 5% and 1% level respectively, indicating rejection of the null hypothesis.

Figure 1
Volatility Estimation Periods

