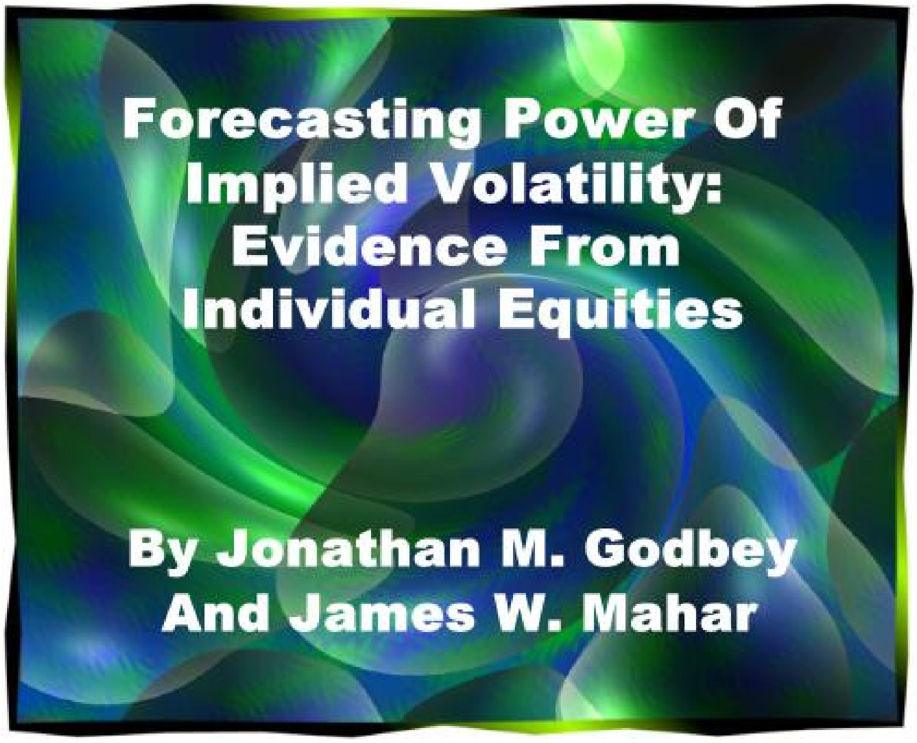

Peer Reviewed

Jonathan M. Godbey godbeyjm@jmu.edu is an Assistant Professor, Department of Finance and Business Law at James Madison University and James W. Mahar jmahar@sbu.edu is an Assistant Professor, Department of Finance at St. Bonaventure University.

Abstract

Assuming use of the correct option pricing model and an efficient market, an option's implied volatility is the market's consensus forecast of future realized volatility over the remaining life of that option. In this paper the authors examine 460 of the S&P 500 firms to demonstrate that: (1) implied volatility is a better forecaster of realized volatility than historic volatility or GARCH models and (2) the information content of implied volatility significantly decreases with liquidity. Since individual equity options are American style, implied volatility estimates were obtained from calls and puts separately, rather than only from calls or pooled data.

Introduction

An accurate forecast of unobservable volatility is a necessary component of virtually all option pricing methodologies. Therefore, rational option market participants will seek the best possible forecast of future realized volatility (RV) over the life of the option. They will estimate RV from both public and private information. Market prices are then set according to an option pricing model, the observed parameters of that model and a volatility estimate that reflects the aggregated forecast of market participants. By setting the result of a particular option pricing model equal to the market price the implied volatility (IV) of that model is obtained. Merton (1973) shows that if there is efficiency in the options market and participants use the correct model, then IV should be the best estimate of RV.

While the widely accepted hypothesis that IV is the market's best forecast of RV is theoretically appealing, empirical testing has proven to be difficult. The joint hypothesis of market efficiency and a correct option-pricing model along with other complexities prevent empirical tests of the information content of IV from being conclusive. Nevertheless, the empirical tests are not without merit. With the unique circumstances of each study in mind, one can develop a better understanding of the information content of IV as well as that of historic volatility (HV) and GARCH forecasts (GAR).

Previous research into the predictive ability of IV for individual stocks used much smaller samples. Lamourex and Lastrapes (1993) examine 10 firms from April 1982 to March 1984. They extract an IV from the Hull and White (1987) model and show that it is biased but a better predictor of RV than HV or GARCH. They further find that HV improves the forecast of IV alone. Mayhew and Stivers (2003) examine 50 firms from 1988 to 1995. IV is obtained using a binomial tree procedure described by Whaley (1993). They find that IV "reliably outperforms GARCH and subsumes all information in return shocks beyond the first lag." They show that for lower volume options IV loses some of its predictive power and may be inferior to approaches relying only on past data.

We use IV estimates from 460 individual stocks from the S&P 500 from October 1, 2001 to September 13, 2002 to test the ability of IV to predict RV. We then compare the predictive power of IV to that of HV and GAR. Because previous research suggests that volume is an important factor in the information content of IV, we rank firms on their option volumes and test the relative predictive power of IV, HV and GAR across volume quintiles. Unlike previous research, the size of our sample is large enough to allow us to conclude that the results are not limited to selected stocks.

We find that IV, HV and GAR all provide useful information in forecasting realized volatility. However, the information content of the variables is not the same. Both alone and in more inclusive models, IV has more predictive ability than HV or GAR. Moreover, when all variables are examined simultaneously, only IV is significant for all quintiles. These findings are consistent with the hypothesized view that IV has higher information content than HV or GAR. Equally interestingly, we show that while implied volatility is a significant predictor of realized volatility, its predictive ability increases with option market volume. This last finding may have important implications for options that trade infrequently.

Our work is unique in that we look at implied volatility from calls and puts separately while previous research relied solely on calls or an index of both. Because individual equity options are American options, the implied volatility based on puts and calls may be different. By keeping the data separate, we are able to show that the implied volatility estimate from each type of option has the similar amount of information and this information content is positively related to the liquidity of the options.

In the next section of this paper we provide a review of relevant literature. Following it is a section which describes the data and tests the methodology. It is followed by a section which contains the empirical findings and a section which concludes the paper.

Literature Review

Poon and Granger (2003) provide an extensive review of the literature related to forecasting volatility. They divide the existing research into two general categories: (1) papers using historical data only and (2) papers using IV alone or in addition to historical data. In general, the latter studies have found that IV contains a significant amount of information and that it is often superior to models that rely on historical information alone.

Since it is reasonable to assume that different markets have differing degrees of efficiency, the forecasting power of IV for one asset class does not necessarily mean that IV will have equivalent capabilities in another. While the testing methodologies may be similar, the results of the IV tests should be considered according to asset class. Following Poon and Granger (2003) these classes are: individual stocks, market indices, currencies, and other assets.

Latane and Rendleman (1976) use weekly data from options on 24 individual stocks from October 1973 to June 1974 to find that IV provides a more accurate forecast than historic volatility (HV). Chiras and Manaster (1978), Schmalensee and Trippi (1978), and Beckers (1981) support Latane and Rendleman's findings for individual stocks. All of these studies use the Black and Scholes (1973) model to find IV. These studies are not conclusive for several reasons: the options market underwent a structural change after the crash of October 1987 (see Rubinstein (1994)); they assume constant volatility; they do not adjust for dividends, or they do not adjust for the possibility of early exercise; and with the exception of Beckers (1981), these studies did not control for the nonsynchronous data.

Lamoureux and Lastrapes (1993) and Mayhew and Stivers (2003) are the only major studies that we are aware of to examine IV's predictive power for individual stocks when compared to conditional heteroskedasticity models. Both studies find that IV is a better predictor of RV than HV or GAR. Mayhew and Stivers (2003) provide the strongest support for IV. They show that IV "captures most or all of the relevant information in past return shocks, at least for stocks with actively traded options." They further show that the predictive power of IV deteriorates with option volume.

Donaldson and Kamstra (2005) compare the predictive ability of implied volatility extracted from call options on the S&P 500 to ARCH models. They find that for high volume periods IV is more informative than ARCH models, but that ARCH models are more informative in low volume periods. Regardless of option volume, they find that IV and ARCH models provide additional information.

In this paper we examine the information content of IV, GARCH based forecasts, and HV for individual stocks. Our sample is significantly larger then previous studies. We examine approximately 460 firms from the S&P 500. Additionally, we examine the differences in the information content of implied volatilities across both puts and calls. To the best of our knowledge, this has not been done in the academic literature. It is important not only because we are dealing with American options, but also due to the previous findings of Mayhew and Stivers (2003) and Donaldson and Kamstra (2005) that report that liquidity is an important determinant in the information content of the option. Given that there may be liquidity differences across markets, puts are analyzed separately.

Data and Methodology

Daily IV for both calls and puts from October 1, 2001 to September 13, 2002 are obtained from iVolatility.com for 460 of the 500 firms in the S&P 500 as of October 1, 2001. IV is computed using the Black-Scholes model and four at-the-money options of differing expirations. The results are then normalized to give estimates of IV for an option with 30 calendar days to maturity. Adjustments are made for dividends and "a proprietary weighting technique factoring the delta and vega of each option participating" in the calculations.

Dividend adjusted stock prices are obtained from by CRSP. To approximate 30 calendar days, prices for 22 consecutive trading days are used to find daily RV and HV. The calculation for annualized volatility of daily returns is as follows:

$$\sigma = \sqrt{252}\left[\sum\left(r_t - \bar{r}\right)^2 / (N-1)\right]^{(1/2)} \tag{1}$$

where $r_t = \ln(S_t/S_{t-1})$ and $\bar{r}$ is the mean return. S_t is the dividend adjusted stock price on day t. Consistent with iVolatility.com's methodology an average of 252 trading days per year is assumed.

GAR forecasts were obtained from a GARCH(1,1) model as developed by Bollerslev (1986). The equation for GARCH(1,1) is:

$$\sigma_n^2 = \gamma V_L + \alpha u_{n-1}^2 + \beta \sigma_{n-1}^2 \tag{2}$$

where V_L is the long-run average variance rate, $u_t = \ln(S_t/S_{t-1})$, σ_n = the estimate of volatility for day n and $\gamma + \alpha + \beta = 1.0$. The maximum likelihood method was used to estimate γ, α, and β by using all available data in the estimation period. The GAR forecast is given by:

$$E\left[\sigma_{n+k}^2\right] = V_L + (\alpha + \beta)^k \left(\sigma_n^2 - V_L\right) \tag{3}$$

Since the parameters are estimated using the entire data set (as in Jorion (1995) and Szakmary et. al. (2003)) the GAR forecasts benefit from information that would not have been available. Hence, our GAR forecasts have a slight advantage over those employed in practice.

We test the hypothesis that (1) IV is an unbiased and efficient estimator of RV, (2) HV is an unbiased and efficient estimator of RV, (3) GAR is an unbiased and efficient estimator of RV, and (4) IV contains all information present in HV and GAR. Our tests follow the methodology of Canina and Figlewski (1993) and Szakmary, Ors, Kim and Davidson (2002). Specifically, we estimate the following regressions:

$$RV_t = \alpha1 + \beta1 * IV_t + \varepsilon_t \tag{4}$$

$$RV_t = \alpha2 + \beta2 * HV_t + \varepsilon_t \tag{5}$$

$$RV_t = \alpha3 + \beta3 * GAR_t + \varepsilon_t \tag{6}$$

$$RV_t = \alpha + \beta1 * IV_t + \beta2 * HV_t + \varepsilon_t \tag{7}$$

$$RV_t = \alpha + \beta1 * IV_t + \beta3 * GAR_t + \varepsilon_t \tag{8}$$

$$RV_t = \alpha + \beta1 * IV_t + \beta2 * HV_t + \beta3 * GAR_t + \varepsilon_t \tag{9}$$

If IV is the unbiased, efficient estimator that theory predicts then α1 and β1 in equation (4) will be zero and one, respectively. If HV is the unbiased, efficient estimator that theory predicts then α2 and β2 in equation (5) will be zero and one, respectively. If GAR is the unbiased, efficient estimator that theory predicts then α3 and β3 in equation (6) will be zero and one, respectively. If IV contains all information present in HV and GAR, then the coefficients of HV and GAR should be zero in equations (7) - (9), while β1 should be significantly different from zero.

To examine the effects of liquidity on the performance of IV, HV and GAR, the stocks are ranked by option volume and grouped into quintiles. This was done separately for call option volume and put option volume. Greater liquidity may be important for a two different reasons. First, it suggests a more cognitively diverse group of investors for a particular stock option. Each investor will bring private information and error into her estimate of the option's price and the resulting IV. Assuming uncorrelated errors, a greater number of investors suggest more information and a better forecast of RV. Second, institutional investors are able to trade more readily in liquid stocks. To the degree that institutional investors may be more rational or have better information than individual investors, it can be expected that IV is a better forecaster in liquid markets.

Unit Root tests (the augmented Dickey-Fuller method) were performed to test the IV time series for stationarity. In all cases the null hypothesis of a unit root could be rejected. Therefore, the sample is assumed to be stationary over the time period studied which allows the use OLS regression for our analysis.

Empirical Findings

We find that (1) IV is a better predictor of RV than HV or GAR, (2) the information content of IV decreases with liquidity of the options, (3) HV and GAR contain significant information about RV not included in IV, (4) neither IV, HV nor GAR appear to be unbiased and efficient estimators of RV, and (5) IV derived from calls and puts produce near identical results. For ease of discussion, we focus on results from call IV. Corresponding and near identical results for put IV are reported in the associated tables.

Tables 1 – 3 (below) summarize the results from the regressions in equations (4) – (6). Regressions are run for each individual stock. The results are then grouped by option volume quintile. Median values for the regression coefficients are reported along with the percentage of stocks where each coefficient is statistically significant at the 5 percent level.

Table 1
$RV_t = \alpha 1 + \beta 1 * IV_t + \varepsilon_t$

Panel A: Call Option Data

Volume	α1	t-value	β1	t-value	r^2	α1%	β1%
1 high	0.102	2.617	0.743	8.180	0.245	0.674	0.946
2	0.107	3.009	0.647	7.713	0.220	0.723	0.936
3	0.122	3.625	0.607	6.735	0.178	0.722	0.867
4	0.095	3.176	0.625	5.952	0.145	0.659	0.868
5 low	0.165	4.320	0.426	4.700	0.095	0.837	0.739
All Firms	0.120	3.455	0.614	6.689	0.176	0.723	0.871

Panel B: Put Option Data

Volume	α1	t-value	β1	t-value	r^2	α1%	β1%
1 high	0.094	2.812	0.747	8.574	0.259	0.546	0.938
2	0.084	2.553	0.663	7.681	0.218	0.593	0.934
3	0.127	3.550	0.596	7.034	0.189	0.652	0.876
4	0.134	3.687	0.548	5.735	0.133	0.692	0.846
5 low	0.141	4.190	0.482	4.711	0.098	0.753	0.774
All Firms	0.115	3.336	0.617	6.721	0.176	0.723	0.875

RV is regressed on IV for each individual firm. The firms are then grouped into quintiles based on the volume of the options used to calculate IV. The median values of the regression results for each quintile are reported above. The percentage of firms with α1 and β1 significant at the 5 percent level are given in columns 7 and 8, respectively.

Table 2
$RV_t = \alpha 2 + \beta 2 * HV_t + \varepsilon_t$

Panel A: Call Option Data

Volume	α2	t-value	β2	t-value	r^2	α2%	β2%
1 high	0.309	11.114	0.237	3.430	0.066	1.000	0.694
2	0.303	11.174	0.209	3.049	0.052	1.000	0.618
3	0.295	11.380	0.158	2.274	0.047	0.989	0.553
4	0.241	10.696	0.249	3.747	0.059	1.000	0.650
5 low	0.241	10.354	0.247	3.861	0.069	1.000	0.726
All Firms	0.275	11.004	0.231	3.369	0.057	1.000	0.647

Panel B: Put Option Data

Volume	α2	t-value	β2	t-value	r^2	α2%	β2%
1 high	0.300	11.020	0.230	3.340	0.059	1.000	0.660
2	0.285	11.090	0.218	3.211	0.064	0.989	0.626
3	0.314	11.353	0.172	2.847	0.047	0.989	0.573

4	0.247	10.801	0.223	3.319	0.055	1.000	0.659
5 low	0.226	10.394	0.248	3.675	0.064	1.000	0.720
All Firms	0.275	10.975	0.231	3.353	0.057	0.999	0.649

RV is regressed on HV for each individual firm. Consistent with Table 1, the firms are then grouped into quintiles based on the volume of the options used to calculate IV. The median values of the regression results for each quintile are reported above. The percentage of firms with α2 and β2 significant at the 5 percent level are given in columns 7 and 8, respectively.

Table 3
$RV_t = \alpha 3 + \beta 3 * GAR + \varepsilon_t$

Panel A: Call Option Data

Volume	α3	t-value	β3	t-value	r^2	α3%	β3%
1 high	0.017	0.207	0.833	3.330	0.052	0.587	0.701
2	-0.006	0.186	0.923	3.247	0.057	0.511	0.630
3	0.010	0.231	0.828	2.838	0.037	0.511	0.620
4	-0.053	-0.708	1.162	3.547	0.059	0.554	0.761
5 low	0.030	-0.455	1.001	3.615	0.057	0.587	0.728
All Firms	-0.028	-0.139	0.967	3.289	0.055	0.546	0.667

Panel B: Put Option Data

Volume	α3	t-value	β3	t-value	r^2	α3%	β3%
1 high	0.024	0.309	0.828	3.248	0.052	0.592	0.684
2	-0.009	-0.089	0.927	3.428	0.100	0.516	0.630
3	0.010	0.203	0.827	2.957	0.071	0.517	0.629
4	-0.053	-0.745	1.160	3.547	0.086	0.560	0.747
5 low	-0.029	-0.452	0.986	3.581	0.056	0.598	0.728
All Firms	-0.017	-0.138	0.940	3.263	0.055	0.558	0.685

RV is regressed on GAR for each individual firm. Consistent with Table 1, the firms are then grouped into quintiles based on the volume of the options used to calculate IV. The median values of the regression results for each quintile are reported above. The percentage of firms with α3 and β3 significant at the 5 percent level are given in columns 7 and 8, respectively.

If IV is an efficient and unbiased estimator of RV then α1 and β1 in equation (2) will be zero and one, respectively. As Table 1 indicates, this is not the case. The median α1 is 0.120, while the median β1 is 0.614. For 72.3 percent of the firms α1 is significantly different from zero, while β1 is statistically significant in nearly 90 percent of the stocks. Panel B of Table 1 shows that the information content of IV is essentially the same for both puts and calls. In each case the median r^2 is 0.176.

Our next finding is that HV is not an efficient and unbiased estimator of RV. As shown in Table 2, the median α2 is 0.275 while the median β2 is 0.231. We find α2 is different from zero for all firms, while β2 is statistically significant for approximately 65 percent of the stocks. HV explains only 5.7 percent of variation.

Results for GAR regression suggest that GAR is less biased than our other variables as estimator. Table 3 shows that the median α3 is -0.028 and is significantly different from zero for only 55 percent of firms. A median β3 of 0.451 suggests that it is not an efficient estimator and is statistically significant for approximately 68 percent of stocks. Similar to HV, GAR explains only 5.5 percent of variation.

As is mentioned above, we separate our sample into quintiles based on the volume of option trading to examine the importance of volume on information content. As Figure 1 (below) shows, the predictive ability (the information content) of IV is positively related to option volume. However, we do not find such a relationship for HV or GAR across volume quintiles. The median r^2 of the IV regression falls with every quintile as do the percentage of firms where β1 is significant. As shown in Tables 2 and 3, liquidity has no discernable effect on the information content for HV or GAR.

Figure I

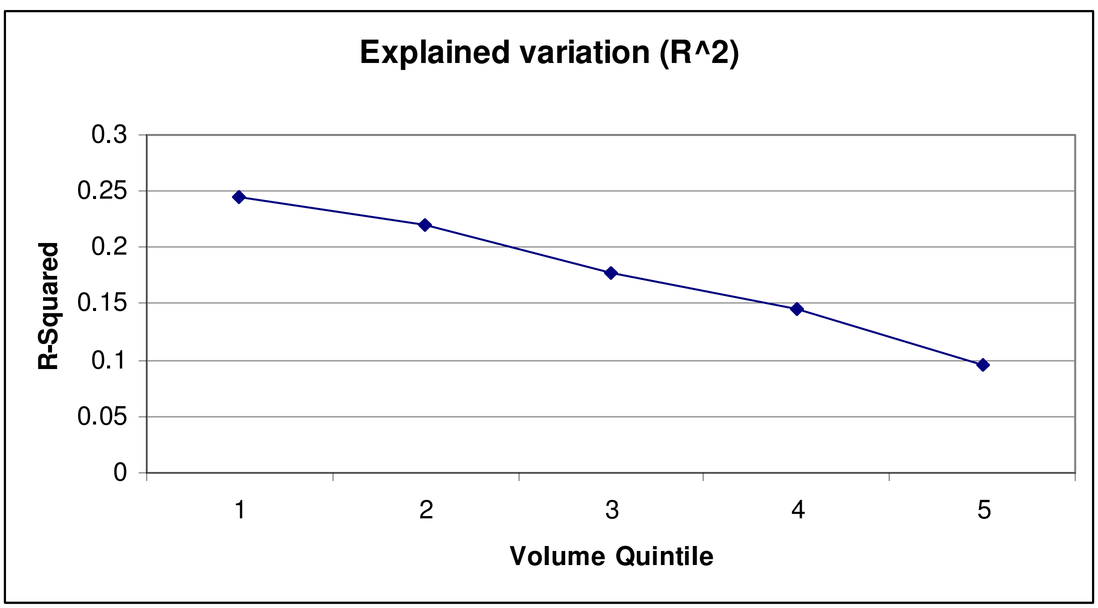

This relation between IV and liquidity is consistent with the findings of Mayhew and Stivers (2003), who report a degradation of IV's information content based on firm size within the 50 stocks that they examined. However, it is important to note that unlike Mayhew and Stivers, we use actual volume numbers, rather than size as a proxy for market liquidity. That information content drops with volume is also consistent with research by Donaldson and Kamstra (2005), who report that the information content of index options is reduced in times of low volume.

Tables 4 and 5 (below) report the results from the multiple regressions in equations (7) and (8). IV is shown to subsume some of the information content of HV and GAR. However, as volume decreases the contribution of HV and GAR increases. Again, this result is consistent with the findings of Donaldson and Kamstra (2005). As Table 4 (below) shows, when RV is regressed on IV and HV, β2 is significant for 21.4 percent of the firms, as opposed to 64.7 percent when IV is not included. For quintile 1, β2 is significant for only 13.0 percent of the firms. This percentage increases steadily across quintiles to 34.8 percent for quintile 5. Table 5 (below) shows similar results for β3. For quintile 1, β3 is significant for only 13.7 percent of the firms.This percentage increases to 38.3 percent for quintile 5. Consistent with Table 1, Table 4 and Table 5 show that the percentage of firms with β1 significant decreases across volume quintiles.

Table 4

$$RV_t = \alpha + \beta 1 * IV_t + \beta 2 * HV_t + \varepsilon_t$$

Panel A: Call Option Data

Volume	α	t-value	β1	t-value	β2	t-value	r^2	α%	β1%	β2%
1 high	0.099	2.855	0.824	7.653	-0.134	-1.878	0.317	0.674	0.924	0.130
2	0.097	3.065	0.877	7.885	-0.217	-3.094	0.305	0.713	0.883	0.128
3	0.129	3.516	0.713	6.022	-0.157	-2.157	0.243	0.722	0.800	0.178
4	0.084	2.831	0.661	4.606	-0.009	-0.115	0.199	0.637	0.780	0.286
5 low	0.164	4.620	0.460	2.859	0.042	0.460	0.157	0.815	0.598	0.348
All Firms	0.128	3.350	0.740	6.007	-0.112	-1.482	0.240	0.712	0.797	0.214

Panel B: Put Option Data

Volume	α	t-value	β1	t-value	β2	t-value	r^2	α%	β1%	β2%
1 high	0.088	2.814	0.855	8.155	-0.174	-2.223	0.336	0.688	0.896	0.149
2	0.069	2.164	0.891	7.400	-0.211	-2.252	0.308	0.637	0.902	0.154
3	0.126	3.468	0.722	6.535	-0.225	-2.983	0.235	0.719	0.820	0.169
4	0.135	3.491	0.615	4.562	-0.022	-0.202	0.178	0.747	0.725	0.253
5 low	0.161	4.226	0.461	2.989	0.084	1.026	0.172	0.787	0.606	0.383
All Firms	0.119	3.157	0.734	6.113	-0.130	-1.722	0.240	0.716	0.790	0.221

RV is regressed on IV and HV for each individual firm. Consistent with Table 1, the firms are then grouped into quintiles based on the volume of the options used to calculate IV. The median values of the regression results for each quintile are reported above. The percentage of firms with α, β1 and β2 significant at the 5 percent level are given in columns 9-11.

Table 5
$RV_t = \alpha + \beta1*IV_t + \beta3*GAR_t + \varepsilon_t$

Panel A: Call Option Data

Volume	α	t-value	β1	t-value	β3	t-value	r^2	α%	β1%	β3%
1 high	0.194	1.774	0.828	7.571	-0.436	-1.763	0.316	0.674	0.916	0.137
2	0.240	1.988	0.765	6.469	-0.403	-1.519	0.267	0.634	0.903	0.129
3	0.225	1.847	0.677	5.985	-0.353	-1.360	0.211	0.544	0.800	0.200
4	0.073	0.890	0.518	4.284	0.150	0.610	0.176	0.438	0.798	0.200
5 low	0.062	0.977	0.431	3.528	0.227	0.905	0.156	0.500	0.564	0.383
All Firms	0.146	1.445	0.665	5.705	-0.233	-0.663	0.214	0.559	0.796	0.228

Panel B: Put Option Data

Volume	α	t-value	β1	t-value	β3	t-value	r^2	α%	β1%	β3%
1 high	0.210	1.972	0.817	7.705	-0.483	-1.825	0.301	0.327	0.908	0.173
2	0.230	2.026	0.783	6.723	-0.393	-1.556	0.275	0.373	0.923	0.165
3	0.164	1.813	0.650	5.613	-0.375	-1.345	0.192	0.416	0.809	0.169
4	0.068	0.793	0.494	4.167	0.184	0.762	0.179	0.538	0.703	0.307
5 low	0.058	0.867	0.380	2.950	0.265	0.915	0.154	.0489	0.598	0.391
All Firms	0.149	1.530	0.665	5.750	-0.221	-0.743	0.212	0.429	0.788	0.241

RV is regressed on IV and HV for each individual firm. Consistent with Table 1, the firms are then grouped into quintiles based on the volume of the options used to calculate IV. The median values of the regression results for each quintile are reported above. The percentage of firms with α, β1 and β3 significant at the 5 percent level are given in columns 9 - 11.

Table 6 (below) shows for approximately 77 percent of the firms, r^2 is higher for IV as the independent variable than it is for HV as the independent variable. Likewise, as shown in Table 7 (below), r^2 is higher for IV as the independent variable than it is for GAR as the independent variable in about 78 percent of the overall sample. This result is consistent with the hypothesis that IV contains more information than just the historical volatility or GARCH models and is consistent with previous research.

Table 6

Predictive Power of IV versus HV

Volume	Calls	Puts
1 high	89.13%	88.86%
2	85.87%	85.71%
3	79.35%	79.77%
4	69.57%	72.53%
5 low	61.96%	59.13%
Total	77.17%	77.22%

The number and percentage of firms by option volume quintile (n=92) where IV offers superior information (as defined as having a higher adjusted r-squared) in forecasting RV. When sorted by volume, 13 firms were dropped due to a lack of volume data. The last row includes the dropped firms.

Table 7
Predictive Power of IV versus GAR

Volume	Calls	Puts
1 high	89.1%	93.5%
2	89.1%	91.3%
3	80.4%	79.3%
4	76.1%	76.1%
5 low	58.7%	60.9%
Total	78.5%	80.2%

The number and percentage of firms by option volume quintile (n=92), where IV offers superior information (as defined as having a higher adjusted r-squared) in forecasting RV. When sorted by volume, 13 firms were dropped due to a lack of volume data. The last row includes the dropped firms.

Conclusion

We use data from 460 of the S&P 500 firms to show that implied volatility is a better predictor of future realized volatility than is historical volatility or GARCH models. We further show that this predictive ability of the implied volatility decreases as option volumes decreases. This is true for both puts and calls. These results indicate that the market for individual equity options is efficient to some degree. They also show that the market's efficiency tends to decline with volume.

Our results have important implications for event studies which use IV to examine changes in firm-specific risk. First, consistent with theory, IV does have significant information content. Therefore, changes in IV signal changes in the market's perception of changes in future return volatility. Second, changes in IV for stocks with active options give a more complete view of the market's perception of changes in future return volatility. Therefore, event studies on more active stocks should produce more robust results than those on illiquid firms.

References

Beckers, S., 1981, Standard deviations implied in option prices as predictors of future stock price variability, *Journal of Banking and Finance 5,* 297-308.

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

Bollerslev, T., 1986, Generalized Autoregressive Conditional Heteroscedasticity, *Journal of Econometrics*, 307-27.

Canina, L. and S. Figlewski, 1993, The informational content of implied volatility, *The Review of Financial Studies 6:3*, 659-681.

Chiras, D.P., and S. Manaster, 1978, The information content of option prices and a test of market efficiency, *Journal of Financial Economics* 6, 213-234.

Christensen, B.J., C.S. Hansen, 2002, New evidence on the implied realized volatility relation, *European Journal of Finance* 7, 187-205.

Donaldson, R.G. and M. Kamstra, 2005, Volatility forecasts, trading volume and the ARCH vs. option-implied tradeoff, *Journal of Financial Research,* forthcoming.

Jorion, P. 1995, Predicting volatility in the foreign exchange market, *Journal of Finance* 1, 507-528.

Lamoureux, C. and W. Lastrapes, 1993, Forecasting stock-return variance: toward an understanding of stochastic implied volatilities, *The Review of Financial Studies 6*, 293-326.

Latane, H.A., and R.J. Rendleman, Jr., 1976, Standard deviations of stock price ratios implied in option prices, *Journal of Finance 31*, 369-381.

Mayhew, S. and C. Stivers, 2003, Stock return dynamics, option volume, and the information content of implied volatility, *Journal of Futures Markets* 23, 615-646.

Merton, R., 1973, Theory of rational option pricing, *Bell Journal of Economics and Management Science,* 4: 141-183.

Poon, S. and C.W.J. Granger, 2003, Forecasting volatility in financial markets: a review, *Journal of Economic Literature*, 478-539.

Rubenstein, M, 1994, Implied binomial trees, *Journal of Finance*, 771-818.

Schmalensee, R. and R.R. Trippi, 1978, Common stock volatility expectations implied by option premia, *Journal of Finance 33*, 129-147.

Szakmary, A., E. Ors, J.K. Kim and W.N. Davidson III, 2003, The predictive power of implied volatility: evidence from 35 futures markets, *Journal of Banking and Finance 27*, 2151-2175.

Whaley, R.E., 1993, Derivatives on market volatility: hedging tools long overdue, *Journal of Derivatives 1*, 71-84.

Note: The graphic containing the title of this paper was designed by Carole E. Scott.