Stock Return Predictability and the Cross Section of Implied Volatilities
Preliminary and Incomplete

Micah Allred*
UCLA Anderson School of Management,

June 10, 2009

Abstract

Motivated by the theoretic work of Lo and Wang (1995), we examine the impact of cross sectional variation in stock return predictability on the cross section of implied volatilities of options. We find that, as tentatively proposed by Lo and Wang (1995), firms with higher predictability have lower implied volatility, on the order of a statistically significant 0.5% percentage points for each quintile increase in predictability. Returns to selling delta-hedged options suggests that the market may not accurately impute the impact of return predictability, with delta-hedged returns increasing with distance from moneyness and decreasing in predictability, suggesting predictability may be mis-priced.

*This paper has greatly benefited from comments from Joshua Coval, Bob Geske, Mark Grinblatt, Hanno Lustig, and Pierre Olivier-Weill as well as participants at UCLA Anderson's Brown Bag Seminar. All remaining errors are my own.

1 Introduction

A great deal of interest, both academic and practical, has been given to the interaction between option prices and the movements and characteristics of the underlying equity market. Recent work has looked at the accuracy of risk neutral implied distributions in relation to realized returns, showing some mis-pricing of higher order moments Ait-Sahalia, Wang, and Yared (2001). Pan and Poteshman (2006), and Xing, Zhang, and Zhao (2008), among many others, investigate the informational content of stock versus option volume as well as the direction that information flows between these two markets. One area that theory suggests should link options and their underlying stocks, however, has yet to receive much empirical interest, namely, the impact of stock return predictability on the levels of implied volatility. Lo and Wang (1995) consider pricing options in a Black-Scholes framework where the underlying is predictable by replacing the standard assumption of Geometric Brownian motion with an Ornstein-Uhlenbeck process, thus introducing mean reversion into the return process, and concluding that the implied volatility used as an input for the BS model must be different from the the standard deviation of log returns if the data generating process is OU rather than Geometric Brownian. Whether that leads to higher or lower implied vol depends on the exact specification of the drift, but it is shown that under a set of conditions[1], that the implied volatility should be lower. The intuition for this result can be gained by considering a decomposition of the unconditional variance, which we assume is fixed, $Var[r(t)] = Var[E[r(t)|\Omega_t] + E[Var[r(t)|\Omega_t]$, where, again, holding the unconditional variance fixed, an increase in the variability of the conditional mean (i.e. predictability) equates to a decrease in the variance of the residual. It is this decrease in residual variance which we would hope to identify in highly predictable stocks. It should be noted that, in a related paper, Liao and Chen (2006) find that implied volatility should be higher for European options when the drift is charac-

[1]The conditions include fixed unconditional variance, the drift not being a function of log price process, and independence of the Brownian processes driving their model

terized as a $MA(1)$ process. It is the intent of this paper to empirically measure and identify the impact of predictability on option prices, and determine both the magnitude, and direction of its effect.

It hardly needs stating that the literature on the determinants of stock return predictability is vast and active. Some of the major contributions to this area include the identification of the size effect, where small firms tend to outperform large firms Banz (1981), the January effect Reinganum (1983), book-to-market and size factors Fama and French (2001), momentum, which shows the recent winners and losers tend to continue to be winners and losers, respectively Jegadeesh and Titman (1993), and long run reversal De Bondt and Thaler (1985)[2]. The intent of these studies is generally to identify factors or characteristics which explain the cross sectional variation of stock returns. However, for the purposes of this paper, we need to identify which stocks have returns that are more predictable than another's. So, while momentum results would suggest that past winners will outperform past losers, we ask, for example, is the error of our forecasted returns using information about past returns greater for past winners or losers? To this end, at the beginning of each month we sort stocks into quintiles based on their return predictability over the previous year, as measured by their mean-squared forecast errors, and show that this non-parametric measure leads to a strongly significant result indicating that greater predictability leads to lower implied volatility. This result is robust to double and triple sorting on various measures of past volatility in addition to past predictability. The forecasting equation is similar the that used in Vuolteenaho (2002), who utilizes past returns, book to market, leverage, and profitability. In addition, we add measures of turnover, volatility, and the dividend yield. The coefficients from the 24 most recent regressions which do not overlap with our forecast period are averaged, then used to forecast the returns for each stock in our sample. Firms which we are best able to predict, in terms of out of sample mean-squared forecast error, over the previous year, are marked as most predictable.

[2]Also important are Daniel and Titman (1997), Ang, Hodrick, Xing, and Zhang (2006), and Chordia, Subrahmanyam, and Anshuman (2001)

We find that stocks which have been most predictable over the past year continue to be more predictable than the other quintiles over a variety of forecasting horizons ranging from 1 to 12 months. It is this non-parametric quintile ranking which is used as the central variable for our further study. We find that the impact of a 1 quintile increase (decrease) in predictability decreases (increases) implied volatility by 0.5%. This effect is fairly constant over increasing time horizons, which implies a greater price impact for longer dated options, consistent with vega, or an options sensitivity to changes in volatility, increasing with time to maturity.

This paper is related to several recent papers which examine links between options and stock returns. A large portion of this literature is devoted to ascertaining the direction of information flow between the two markets. Xing, Zhang, and Zhao (2008) show that the smirk of individual options is a significant predictor of future returns, with a portfolio long low-smirk stocks and short high-smirk stocks earning a risk adjusted 11% per year. They argue that information, particularly negative information, is first incorporated into option prices. Pan and Poteshman (2006) obtain related results by showing that option volume is predictive of future equity returns, and go further, showing that the predictability is driven by private information (in their case, access to order flow by client type provided their valuable private information)[3]. These results are consistent with Easley, O'Hara, and Srinivas (1998) who derive a model similar to Easley and O'Hara (1987) which shows that informed traders should focus on option markets where available leverage is higher. Chakravarty, Gulen, and Mayhew (2004) investigate how price discovery is split between option and stock markets, and find that informed traders tend to value the leverage provided by options, and that trading in options significantly contributes to stock price discovery. Most closely related to this paper would be a working paper by Seyhun and Wang (2006) which uses short term (5 days) returns and autocorrelations to explain violations to put call value boundaries in the spirit of Amin, Coval, and Seyhun (2004). Their sample is limited to 5 stocks which have a high volume

[3]Other notable contributions are Bollen and Whaley (2004),Dennis and McConnell (1986),Toft and Prucyk (1997),Bakshi and Kapadia (2003), and Longstaff (1995)

of intraday trading which is necessary to match trade times, and while recognizing the theoretic link between return predictability and implied volatility, state explicitly that the investigation of that relationship is not something they intend to pursue. This paper contributes to the extant literature by focusing on the variable theory suggests should be affected by predictability, namely, the implied volatility. The data set covers over 700 firms during the 12 year span from 1996-2007, and utilizes a panel data methodology to account for both time and firm fixed effects, as well as leverage, recent returns, persistence of volatility along with other control variables.

The paper will proceed as follow: section I outlines the data, section II provides the details of the methodology, section III examines the main results, section IV examines the outcome a trading strategy, and section V concludes.

2 Data

Our options data comes from OptionMetrics which provides a variety of characteristics of the option including end of day implied volatility, delta, closing bid and asks, and the maturity of the option. The data cover trading days between January 1996 and December 2007. The moneyness of a given option is determined using the delta of the option, following Bollen and Whaley (2004). As Bollen and Whaley point out, the delta provides a close proxy to the probability that an option will finish in the money [4] and also incorporates the volatility of each firm as well as the time to expiration of the given option, facilitating the comparison between stocks which have different implied volatilities, and options with varied times to maturity. Table 1 shows the delta ranges assigned to each moneyness category. The implied volatility is calculated by OptionMetrics by iterating over values of σ_{impl} within a CRR tree. The data is filtered by first removing observations where the

[4] $N(d_2)$ is the correct probability of the stock finishing in the money under the risk neutral measure, although delta, or $N(d_1) = N[\frac{ln((S-PV(dividends))e^{r}T/X)+\frac{\sigma^2 T}{2}}{\sigma\sqrt{T}}]$, is commonly used in its place

implied volatility is less than 2% or greater than 150%, delta is less than .02 or greater than .98, as well as duplicate observations for the same optionid, which often violate put-call parity[5]. On any given day there may be multiple options, both calls and puts within a given moneyness category, for which we have quotes. All valid observations are averaged to provide one observation per maturity and moneyness bucket. Data on stock returns and corporate characteristics such as book-to-market or leverage, are drawn from CRSP and Compustat.

In Table 2 we have the summary characteristics of the stocks by predictability quintile, both for the entire sample of all stocks from 1970 to 2007, and in Table 3, we present the same summary statistics for those firms which have options in our 1996 to 2007 sample. For the entire sample of CRSP firms we find that are the most predictable are generally larger, with an average size of $\$5.7B$, and average size decreases monotonically with decreasing predictability to $\$1.8B$ for the least predictable. The market cap for the lowest quintile may seem unusually high, however, many lower market value firms are excluded from consideration as a firm needs to have roughly three years of trading (further details will be given in the methodology section) to be given a predictability ranking, thus the firms here are larger than the general universe of CRSP stocks. The average bid-ask spread and turnover increase as the firm size declines, while predictability decreases. [6]. The patterns for book to market and leverage are less dramatic, but are broadly consistent with smaller firms having lower book to market ratios due to their smaller amount of assets and greater value of future growth options, as well as being more funded by equity, leading to lower leverage. For those firms which have options in our sample, we see that the most predictable firms are still the largest, with a market value of approximately $\$11.5B$ compared with $\$4.0B$ for

[5]Xing, Zhang, and Zhao (2008) consider any option which has positive open interest without a restriction on recent trading. The main panel regression results of this paper where both qualitatively and quantitatively similar when replicated with their filter.

[6]It is commonly argued in the market microstructure literature that high trading volume, when generated by symmetric shocks to buy and sell volume, is a signal of a lack of consensus regarding the prospects of the firm Duarte and Young (2008). The increase of turnover associated with the decrease in predictability appears supportive of that conjecture

the lowest, placing them among the largest firms in the full CRSP universe. For the firms that have options included in our sample, we see that, as anticipated, the firms are among the largest in CRSP, and correspondingly, the dispersion in market equity, turnover, and spreads, is much lower than in the broader market.

One may hypothesize that variation in implied volatilities associated with the different levels of predictability is driven by industry effects rather than by predictability. Moskowitz and Grinblatt (2002) find that the momentum returns are largely determined by industry level momentum, and, in our case, perhaps predictable firms are simply all large utilities and our results are driven by industry effects. Looking at Table 4, we can see that 44% of the firm month observations for firms in the utilities sector were in the highest predictability quintile, while only 5.3% are in the lowest. Conversely, industries like Health Care, Computers, and Chips$/$electronics, with a high percentage of firm month observations in the lowest predictability quintile also have a low incidence of high predictability observations. It thus is possible that there may be industry effects which will need to be addressed.

3 Methodology

Our first priority is attempting to identify which stock returns are more predictable than others, that is, which ones have the lowest forecast error variance. Thus, rather than saying, for example, that a portfolio of stocks which have been winners over the previous 6 months will outperform a portfolio of past losers, we wish to say, stocks which we have been able to accurately forecast returns over the past year will continue to be more forecastable than others. This is our link to the predictable drift term in Lo and Wang (1995) or Liao and Chen (2006). While we, along with any other market participant, are obviously unable to identify the exact specification of the drift term, if we are able to identify those stocks whose returns are the most predictable, then market participants would be expected to decrease (or increase) the implied volatility relative to other stocks to account for that variation in predictability. It is the extent to which the market adjusts for predictability, and its accuracy in doing so, that we wish to explore.

To create our return forecast we pull from some of the more common and successful variables suggested by the literature including book to market Fama and French (1993) and Daniel and Titman (1997), past returns Jegadeesh and Titman (1993), profitability Haugen and Baker (1996), turnover Chordia, Subrahmanyam, and Anshuman (2001), volatility Ang, Hodrick, Xing, and Zhang (2006), dividend yield Hodrick (1992) and Campbell and Shiller (1988), and leverage Bhandari (1988). This return forecasting equation is very similar that used by Vuolteenaho (2002) as part of his VAR based decomposition of firm level returns.

$$r_{i,t} = \Sigma_{i=12}^{1} r_{t-i} + D/P_{t-1} + Turn_{t-1} + Vol_{t-1} + Profit_{t-1} + BM_{t-1} + Lev_{t-1}$$

To obtain the coefficients we will use in the forecast, we start by regressing the above list of variables at the beginning of each month on future 1,3,6, and 12 month returns (separately) , and the coefficients are stored. Once 24 months of stored coefficients are available, they are averaged and used to create an out of sample return forecast. The stocks are sorted on their mean-squared forecast error over the previous 12 months, or rather, over the 12 investment horizons that terminated within the last year, since in the case of 3 month horizon returns, the first 3 month period ending in a given year would also include returns for November and December of the previous year. Returns are then forecasted over the next 1,2,6, or 12 months and the mean-squared forecast error is calculated for each quintile. This process is repeated each month for the entire sample, covering the period 1970 - 2007. As the first 24 months are used to generate the coefficients for the forecast equation, and (at least) the next 12 months are needed to create the MSE over which the firms are sorted into quintiles, the sample is reduced to 1974 through 2007. Table 5 shows the average mean-squared forecast errors by quintile. The MSE declines monotonically from most predictable to least, indicating that those firms which have been more forecastable, will continue to be predictable. Although not presented here, the differences in MSE between quintiles are statistically significant using the standard errors have been calculated with the Newey-West procedure. Naturally, this is a noisy measure of predictability. As such, we consider two and three-way sorts controlling for past one month volatility and past one year (excluding the past one month) volatility Table 6 and Table 7. In the Two Way sort, we sort on the past one month volatility and the MSE rank, and still find that even within the same recently volatility quintile, the mean-squared pricing error decreases as we move from lower predictability quintiles to higher ones. Additionally, in Table 7, we first sort on past one year, and then past one month volatility, and find that the same pattern generally persists, where increases in past predictability rankings lead to lower mean squared forecast errors. This exercise will hopefully allay concerns that our measure of predictability is simply a proxy for volatility.

To test the impact of return predictability on option implied volatility, we need to first create a cross section of option implied volatilities. All of these observations are averaged each calendar week within each maturity and moneyness category. Moneyness categories range from 1 through 5 as seen in Table 1, where category 1 represents deep in-the-money (DITM) calls and deep out-of-the-money (DOTM) puts, 3 represents at-the-money (ATM) calls and puts, and 5 is for DITM puts and DOTM calls. Maturity buckets are arranged such that the most liquid options, those expiring in the next two months, are in bucket 1, bucket 2 is for those expiring between 2 and 4 months out, 3 is for 4 to 6 months out, 4 is for 6 to 12 months, and 5 for those options expiring over a year from the current date.

4 Results

This set up lends itself nicely to panel data methods for unbalanced panels. As our ranking on predictability is updated monthly, our time series frequency is monthly. We take the average implied volatility within each moneyness and maturity category during the second week of the month, include both firm and time fixed effects, and adjust the standard errors for serial correlation and heteroskedasticity with the panel version of Newey-West errors. Separate panels are calculated for each maturity bucket to address maturity specific effects as well as to match maturity specific predictability ranks with the appropriate maturity of options. That is, we use predictability rankings for 1 and 3 month horizons for our bucket of options that expire within the next 2 months, and 3 and 6 month rankings are used for options with 2 to 4 months to maturity etc. We also choose the two way fixed effects panel approach in order to address concerns about biased standard errors in Fama-MacBeth regressions raised by Peterson (2008). Peterson demonstrates the potentially significant bias generated when there are firm effects and standard errors are generated in a Fama-MacBeth setting, among others. It is noted that Peterson advocates using clustered standard errors, however, clustered errors were also calculated and where essentially the same to those presented, and thus are not included. While our interest is in predictability, we must be careful to ensure that our results are not driven by some other factor. We have already pointed out that the returns of certain industries tend to be more, or less, predictable, and hence, we include as a control variable the previous months realized standard deviation for each industry[7] to address industry specific uncertainty. Other results in the literature have shown that higher firm leverage increases the volatility of its stock returns Toft and Prucyk (1997), Christie (1982). We thus we calculate and include firm leverage. We also include the past 4 weeks return to address the tendency for implied volatility to increase after positive returns and decrease after negative ones Bakshi and Kapadia (2003). It is

[7]This is calculated as the arithmetic average of realized standard deviations for all firms identified as being in the industry

also worthwhile to note here that this variable, past 4 week return, is related to the measure used in Seyhun and Wang (2006), who use past 1 week return and state that their results are robust to using past 2 week returns, to identify violations to put call price boundaries due to buying pressure induced by predictability. Additionally, it is a well known fact that volatility clusters, with periods of high and low volatility grouping together, hence the wide spread use of GARCH/ARCH models Bollerslev (1987), Engle (1982). We therefore include both the past 4 week realized volatility along with the realized volatility over the period 1 month to 12 months previous, which we interpret loosely as the long term variance and the present innovation to the volatility process.

Looking at Table 8, we see that, as expected, firms with higher leverage have higher implied volatilities, and those which have experienced large recent negative returns, hence an increase in leverage, also have higher implied vol indicated by the significant negative coefficient of $-.124$ for the return variable in the one month regression. Consistent with the expectation of persistence in volatility, the coefficients of both the past 1 month, and 1 year realized volatility are both extremely significant, with coefficients of $.19$ for the one month volatility and $.46$ for the one year volatility. The recent industry volatility is similarly positive and significant, though smaller in magnitude and significance to the own firm volatility, and does not remove the effects of our predictability quintile variable, rank. Rank, which here is the rank for 1 or 3 month predictability[8] has a coefficient of 0.43%, indicating that for every quintile increase, and higher quintiles are associated with higher MSE, and thus lower predictability, implied volatility increases by 0.43%. This magnitude of this coefficient has a small range from 0.43% for the one month horizon to $.05\%$ for the 6 month horizon 9, but as an option's vega increases with maturity, the impact of predictability clearly increases with maturity. One alternative explanation for variation in implied volatility is the impact of higher order moments. Under Black-Scholes pricing, the price process is log-normally distributed, or the log price process is normal. In this setup, both of these distributions are fully characterized by

[8]Results are similar for 6 and 12 month sortings and will be examined in more detail in a later table

their first two moments. However, it has long been known, both academically and practically[9], that actual and implied risk-neutral return distributions are non-normal, suffering from skewness and excess kurtosis. It is the non-normality implied by these higher order moments, particularly the skew, that is associated with the smirk of the implied volatility function Bakshi, Nikunj, and Madan (2003), Dennis and Meyhew (2002). To consider the possibility that investors, in addition to bidding up the vol of deep OTM puts for stocks with greater return skewness, also bid up the level of ATM implied volatility, we re-run our regression with return skewness over the past 1 year as a new control variable. To also consider the possible impact of jumps, as described by excess kurtosis, we will also include realized return kurtosis over the past year. The inclusion of skewness produces a coefficient of -0.003 which is statistically significant with a t statistic of -3.4. This is consistent with the notion that negative skewness induces *higher* relative implied volatility, even in ATM options, while those with positively skewed returns merit an implied vol discount. In the far right column, we run the same panel regression, but now include both realized skewness and kurtosis over the past year as explanatory variables. The coefficient for skewness increases slightly to -0.0033, but increases in significance with a t stat of 4.74. The kurtosis variable is $-.001$ with a t statistic of 11.85. Somewhat surprisingly, the sign of the kurtosis coefficient is negative, rather than positive, which would be expected for a stock with more frequent large jumps. Bakshi, Nikunj, and Madan (2003) arrive at a similar result, when using the risk neutral kurtosis implied by options as a regressor in a regression addressing the determinants of delta hedged gains. They posit the existence of some interaction between skewness and kurtosis that generates the negative coefficient. More importantly, even after the addition of higher order moments, the significance of the rank variable is still unchanged, although the magnitude dropped from $.0049$ to $.0043$. As a check to our implicit assumption of a constant linear impact of changes in predictability on implied volatility, we remove the rank variable, and in its place, use 4 dummy variables as seen in Table 10. The increase from one quintile to the next is generally close to 0.5%, validating the use of our

[9] See Taleb (2006) and Heston (1993).

linear rank variable.

Table 11shows the impact of predictability generated implied volatility variation across maturity and moneyness. As we made an effort to match predictability horizons with option maturities previously, here too, we use 0.5% as the per quintile increase broadly consistent with Tables 8 and 9. We can interpret this table under the umbrella of our unconditional variance decomposition. Consider a set of stocks with the same unconditional variance but different levels of predictability. If we imagine the least predictable stock has no variation of the conditional mean, then the unconditional volatility is the same as the expectation of the conditional volatility. If such a stock is in the least predictable quintile, and we move up one quintile in predictability, then the implied volatility of more predictable stock would be, on average, $.04 - .05\%$ less. Similar to Lo and Wang (1995) predictability has minimal impact to short dated options, but the price impact, in absolute terms, increases as maturity increases, although it declines modestly in percentage terms. Also, the percentage impact is much greater for DOTM calls, which fall into our category 5.

5 Trading

In this section we consider the accuracy of the markets pricing of differential predictability. Since we do not necessarily know *a priori* what the correct implied volatility is, we resort to testing the set of trading strategies to see if abnormal returns can be generated by selling and appropriately hedging options, suggesting that options markets mis-price the level of predictability in stock returns.

A number of articles have been written analyzing the profitability of writing options of different moneyness levels and maturities. Whaley (1986) tests american futures options on the S&P 500 in the early 80's, against the price generated by the Black-Scholes partial differential equation with boundary conditions adjusted for American exercise, and finds, in addition to systematic moneyness and maturity biases, that significant risk adjusted profits could be earned by selling OTM puts. He suggests this is evidence of market inefficiency. However, a more recent set of papers have suggested that the returns to writing options, in particular ATM options, are compensation for volatility risk. Bakshi, Nikunj, and Madan (2003) find support for a negative volatility risk premium by looking at the returns to a delta-hedged strategy in S&P 500 index options and find negative returns to buying the option and delta hedging. This is interpreted as paying an insurance premium for the negative correlation between market returns and volatility. Taking a slightly different approach, Bollen and Whaley (2004), argue that excess returns, or conversely, costs, to writing OTM puts for index options, or OTM calls for individual stocks are driven by excess demand for those securities. They argue that if the market demand for a particular option is all on the long side, then the market maker must by net short. To compensate the market maker for taking on the position, a ”volatility markup” is added Green and Figlewski (1999). Consequently, they show that selling deep out of the money puts (and calls), generate 1.8% monthly for highly liquid individual stocks and 2.3% for index options.

We begin here by examining the difference between realized and implied volatilities by predictability quintile. Bollen and Whaley (2004) find that, contrary to index options, the difference between implied and realized volatility for ATM options is negative, indicating that in vol terms, ATM options are cheap, and DOTM and DITM options are somewhat expensive, with implied vol being larger than realized. We find very similar results with a much broader and longer data set. Figure 1 shows the difference between implied and realized volatility by predictability quintile for all options during the first half of the sample period. First, we notice that the shape of the each of the curves are essentially the same, and that the difference curves show a monotonic increase from most predictable on the bottom to least predictable on the top. While the 4 most predictable quintiles are fairly closely grouped together, the least predictable quintile is shifted up, and ATM options have realized vols essentially the same as the implied. At the extreme left of the figure, the

DOTM puts in moneyness category 1 have implied vols that, on average, are close to 8% higher than what is actually realized. This pattern suggests that ATM options for high predictability firms may be underpriced, particularly when transaction costs or vega hedging costs are included, while those for OTM puts may by over priced. Alternatively, it may be that OTM puts for low predictability firms have exposure to higher order moments which may create severe negative skewness in the return distribution or are priced to reflect significant demand. These alternatives will be addressed in the following section.

5.1 Delta Hedged Returns

To generate our returns, we implement a procedure similar to Bollen and Whaley (2004), where we sell each option which is available during the fourth week of each month, on the day of the week closest to Wednesday so as to assuage concerns regarding Monday and Friday trading Xing, Zhang, and Zhao (2008). The fourth week is chosen as it provides a one month return for options which expire the following month, and hence creates a non-overlapping return series. For longer dated options, this set up allows for more consistent overlaps in time for our observations, which will require the use of a Newey West adjustment to the standard errors. Then, for each week until the expiration week of the option, the position is rehedged on the day closest to Wednesday of each week, including the last week when the options expire on Friday. Returns are then separated by maturity, moneyness, and predictability quintile. Following Bollen and Whaley (2004), we calculate the abnormal returns associated with writing an option and delta hedging by holding $N(d_1)$ of the underlying. This position in the underlying may entail the receipt or payment of dividends, which, along with any gains or loses incurred when rebalancing the hedge, are carried forward at the risk free rate until the position is closed out. Explicitly, the abnormal return for a call is calculated as

$$
\begin{aligned}
ARET_c &= \\
&= \frac{\Delta_0(S_T + \sum_{t=0}^{T} D_t e^{r(T-r)} - S_0 e^{rT}) - (c_T - c_0 e^{rT}) + \sum_{t=0}^{T-1} \Delta_t (S_{t+1} + D_t - S_t) e^{r(T-t)}}{\Delta_0 S_0 - c_0}
\end{aligned}
$$

S_t is the closing price from CRSP for date t, Δ_t is the options delta calculated by OptionsMetrics for date t, and c_t is the call options price on date t. The risk free rate r is obtained by interpolating the zero yield curve available in the OptionMetrics database. It is the rate obtained from this curve that is used to move dividends (D_t) and capital gains and losses forward in time to the close out of the position. The three terms in the numerator can be interpreted as, starting from the left hand side, the income generated by holding the initial delta amount of the underlying stock net of financing costs, the gain on the option you sold, hence the negative sign, and finally, the gains or losses due to weekly rebalancing our hedge carried forward to expiration at the risk free rate. In the limit where continuous delta hedging is possible, and volatility risk is not priced, the abnormal returns to this position should be zero. This result can be generalized in that even with discrete hedging, the tracking errors are symmetric around zero such that expectation is still zero Bakshi and Kapadia (2003), Bertsimas, Kogan, and Lo (2000).

As a baseline of comparison, Table 12 shows in the top panel, the abnormal returns to a a strategy similar to the one outlined above, with the major exception that rebalancing was performed daily, for a sample of 20 large firms. This corresponds to the bottom line in Table 8 in Bollen and Whaley (2004). The lower panel displays the results from our strategy of rebalancing weekly, over a similar time period, 1996-2001 rather than 1995-2000 as a result of data limitations. The two rows compare favorably and, with ATM options returning 0.191% in Bollen and Wheley versus 0.10% in our results. Across all categories, the mean abnormal returns exhibit a similar

U shaped pattern moving from DOTM puts to DOTM calls (categories 1 to 5), with the daily re-hedging strategy returning comparable values to the weekly strategy. Although tracking errors are symmetric for relatively small numbers of discrete hedging, our weekly re-hedge is less frequent than commonly considered in Bakshi and Kapadia (2003) or Bertsimas, Kogan, and Lo (2000). However, given the the comparison in Table 12, it appears that any tracking error due to discrete (weekly) hedging will not likely significantly impact our conclusions.

As we look to our trading results, we make one significant change to the Bollen Whaley approach, by considering only OTM and ATM calls and puts, that is, categories 1 and 2 have only puts, and categories 4 and 5 have only calls, while category 3 contains both ATM calls and puts. For much of our broad sample of firms, weekly rehedging for DITM firms, those that have high deltas in absolute terms and tend to be less liquid, led to frequent large negative delta hedged returns. Bakshi and Kapadia (2003) find similar results and attribute the anomaly to stale quotes for the less liquid DITM options, which they also exclude from their analysis. Excluding DITM options also decreases any potential due to jump risk as this is greatest for ITM options Bakshi and Kapadia (2003). As we examine the statistical significance of the delta hedged returns, we note the existence of high skewness in the sample, and hence calculate the Johnson (1978) modified t stat as

$$t_j = (ARET + \sigma S/6n + ARET^2 S/3\sigma)/(\sigma/\sqrt{n})$$

where ARET is the average abnormal return, S is the skewness, n the number of observations, and σ is the standard deviation. With that, we turn our attention to Table 13, where we have the outcome of the delta hedged strategy for the first and last halves of the sample period, as well as that of the entire period. Similar to the results presented earlier, along with those in Bollen and Whaley, we find that ATM money options yield the lowest abnormal returns, which, although statistically significant, are small in magnitude in comparison to the other moneyness categories.

Generally, abnormal returns increase with distance from ATM and with deceasing predictability. We thus find the common U shaped pattern, consistent with the pattern found between implied and realized volatilities, and near monotonic increases in returns going from most to least predictable quintiles. It may be surprising at first to see that the returns to selling delta hedged options which are close to ATM, and hence, generally have higher realized than implied volatility, are positive. This is the result which would be expected if volatility risk were not priced, however, as shown in Bakshi and Kapadia (2003), delta hedged returns need not be zero when volatility risk is priced, and hence enters the option price through the risk neutral drift.

Returns for the most predictable quintile range from 0.4% for ATM options, to approximately 6% for DOTM calls. For the least predictable quintile, categories 1 and 5 see statistically significant returns of 20 to 30%. The numbers here are *monthly* returns, which, given the magnitude of the returns, suggests a potentially fabulous trading opportunity, or more likely, compensation for exposure to volatility risk, transaction costs, and supply and demand imbalances Bollen and Whaley (2004). Table 14 shows the returns when the option is sold at the bid and bought back at the offer, and the hedging is similarly executed at the bid and ask prices rather than the midpoint. Incorporating the bid-ask spread reduces returns across the board, and leads to negative returns for ATM options. All returns are significantly different from zero with the exception of category 2, rank 0. In both the mipoint and bid-ask spread tables, the difference in ATM returns between the least predictable and most predictable is 2.1% and 1.2%, which brackets our estimate of the spread generated by predictability in our panel regressions of 1.6%. Similar results (not shown) hold for panel regressions on OTM options, so only part of the very large spread between category 1 and 5 options is likely attributable to variation in predictability.

One potential explanation for the high delta hedged returns is compensation for providing protection against volatility risk. To consider this, we follow Bollen and Whaley (2004) and consider a strategy similar to the delta hedged one,where during the fourth week of each month, options are sold, and then vega hedged with the option closest to ATM (which has the highest vega), and expires after the option being sold. The net delta to these two options is then delta hedged. Each week the vega hedge is adjusted, and then the net delta is set to zero by trading in the underlying. All transactions are executed at the midpoint. Bollen and Whaley find that abnormal returns disappear when transaction and vega hedging costs are considered, suggesting that if anything, market makers may underprice volatility risk. Our results, shown in Table 15, show that returns are very skewed, yet still often positive, though greatly reduced from the returns to the delta hedged strategy.

6 Conclusion

Considering the theoretical motivation that the existence of predictability in the data generating process can lead to implied volatilities that are potentially either greater than, or less than, those estimated for the standard Black-Scholes model, we attempt to empirically measure the extent to which variations in predictability are incorporated into option prices. Using a panel data set over the period 1996 through 2007, we find that for ATM options with less than 60 days until expiration, a one quintile increase in predictability leads to 0.43% decrease in implied volatility after accounting for persistence in volatility, recent stock price movements, industry effects, leverage, and higher order moments of the underlying. Matching option maturities with return predictability rankings for that horizon, we find that the impact of predictability ranges from 0.43% for the nearest (1 month), 0.48% longest (12 months) horizon, and 0.5% for the intermediate (3 and 6 month) maturities.

Abnormal returns to selling a delta-hedged option are significantly positive when transacted at the midpoint, but are greatly reduced when the trader must trade against the bid-ask spread. In both strategies, the spread in returns for ATM options is close to the value estimated in our panel regression attributable to variation in predictability. It is, however, surprising that although the mean returns to this strategy have declined over time, the returns to selling options is sizable for OTM categories. In this paper we have focused on the return predictability of stocks impacting options, but the basic total variance decomposition works in the other direction. One avenue for future research to examine the ability of changes in implied volatility to forecast when stock returns will be more predictable.

References

Ait-Sahalia, Yacine, Yubo Wang, and Francis Yared, 2001, Do Option Markets Correctly Price the Probabilities of Movement of the Underlying Asset?, *Journal of Econometrics* 102, 67–110.

Amin, Kaukshik, Joshua Coval, and Nejat Seyhun, 2004, Index Option Prices and Stock Market Momentum, *Journal of Business* pp. 835–873.

Ang, Andrew, Robert Hodrick, Yuhang Xing, and Xiaoyan Zhang, 2006, The Cross Section of Volatility and Expected Returns, *Journal of Finance* 61, 259–299.

Bakshi, G., and N. Kapadia, 2003, Delta-Hedged Gains and the Negative Volatility Risk Premium, *Review of Financial Studies* 16, 527–566.

Bakshi, Gurdip, Kapadia Nikunj, and Dilip Madan, 2003, Stock Return Characteristics, Skew Laws, and Differential Pricing of Individual Equity Options, *Review of Financial Studies* 16, 101–143.

Banz, Rolf, 1981, The Relationship Between Return and Market Value of Common Stocks, *Journal of Financial Economics* 9, 3–18.

Bertsimas, D., L. Kogan, and A. Lo, 2000, When is time continuous?, *Journal of Financial Economics* 55, 173–204.

Bhandari, Laxmi, 1988, Debt/Equity Ratio and Expected Common Stock Returns: Empirical Evidence, *Journal of Finance* 43, 507–528.

Bollen, N., and R. Whaley, 2004, Does Net Buying Pressure Affect the Shape of Implied Volatility Functions, *Journal of Finance* 59, 711–754.

Bollerslev, Tim, 1987, Generalized Autoregressive Conditional Heteroskedasticity, *Journal of Econometrics* pp. 307–327.

Campbell, John Y., and Robert J. Shiller, 1988, The Dividend-Price Ratio and Expectation of Future Dividends and Discount Factors, *Review of Financial Studies* 1, 195–228.

Chakravarty, Sugarto, Huseyin Gulen, and Stewart Mayhew, 2004, Informed Trading in Stock and Option Markets, *Journal of Finance* 59, 1235–1257.

Chordia, Tarun, Avanidhar Subrahmanyam, and V. Ravi Anshuman, 2001, Trading Activity and Expected Stock Returns, *Journal of Financial Economics* 59, 3–32.

Christie, Andrew, 1982, The stochastic behavior of common stock variances : Value, leverage and interest rate effects, *Journal of Financial Economics* 10, 407–432.

Daniel, Kent, and Sheridan Titman, 1997, Evidence on the Characteristics of Cross Sectional Variation in Stock Returns, *Journal of Finance* pp. 1–33.

De Bondt, Werner, and Richard Thaler, 1985, Does the Stock Market Overreact?, *Journal of Finance* pp. 793–805.

Dennis, Debra K., and John J. McConnell, 1986, Corporate Mergers and Security Returns, *Journal of Financial Economics* 16, 143–187.

Dennis, Patrick, and John Meyhew, 2002, Risk Neutral Skewness: Evidence from Stock Options, *Journal of Financial and Quantitative Analysis* 16, 143–187.

Duarte, Jefferson, and Lance Young, 2008, Why is PIN Priced?, *Journal of Financial Economics* Forthcoming.

Easley, David, and Maureen O'Hara, 1987, Price, Trade Size, and Information in Securities Markets, *Journal of Financial Economics* 19, 69–90.

——— , and P.S. Srinivas, 1998, Option Volume and Stock Prices: Evidence on Where Informed Traders Trade, *Journal of Finance* 53, 431–465.

Engle, Robert, 1982, Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation, *Econometrica* 50, 987–1008.

Fama, Eugene, and Kenneth French, 1997, Industry Costs of Equity, *Journal of Financial Economics* pp. 153–193.

Fama, E., and K. French, 2001, Disappearing Dividends: Changing Firm Characteristics or Lower Propensity to Pay, *Journal of Financial Economics* 60, 3–43.

Fama, Eugene F., and Kenneth R. French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.

Green, Clifton, and Stephen Figlewski, 1999, Market risk and model risk for a financial institution writing options, *Journal of Finance* 54, 1465–1499.

Haugen, Robert, and Nardin Baker, 1996, Commonality in the Determinants of Expected Stock Returns, *Journal of Financial Economics* 41, 401–439.

Heston, Steven, 1993, A Closed Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options, *Review of Financial Studies* 6, 327–343.

Hodrick, Robert J., 1992, Dividend yields and expected stock returns: Alternative procedures for inference and measurement, *Review of Financial Studies* 5, 257–286.

Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency, *Journal of Finance* pp. 65–91.

Johnson, Norman, 1978, Modified t tests and confidence intervals for asymmetrical populations, *Journal of the American Statistical Association* 73, 536–544.

Liao, Szu-Lang, and Chao-Chun Chen, 2006, European options when asset returns are autocorrelated, *Journal of Futures Markets* 26, 85–102.

Lo, Andrew, and Jiang Wang, 1995, Implementing Option Pricing Models When Asset Returns are Predictable, *Journal of Finance* 50, 87–129.

Longstaff, Francis, 1995, Option Pricing and the Martingale Restriction, *Review of Financial Studies* 8, 1091–1124.

Moskowitz, Tobias, and Mark Grinblatt, 2002, Do Industries Explain Momentum, *Journal of Finance* pp. 1249–1290.

Pan, Jun, and Allen M. Poteshman, 2006, The Information in Option Volume for Future Stock Prices, *Review of Financial Studies* 19, 871–908.

Peterson, Mitchell A., 2008, Estimating Standard Errors in Finance Panel Data Sets: Comparing Approaches, *Review of Financial Studies* 22, 435–480.

Reinganum, Marc, 1983, The anomalous stock market behavior of small firms in January : Empirical tests for tax-loss selling effects, *Journal of Financial Economics* 12, 89–104.

Seyhun, Nejat, and Xuewu Wang, 2006, Past Stock Returns and Option Prices, University of Michigan Working paper.

Taleb, Nassim, 2006, *The Black Swan: The Impact of the Highly Improbable*.

Toft, Klaus, and Brian Prucyk, 1997, Options on Leveraged Equity: Theory and Empirical Tests, *Journal of Finance* pp. 1151–1180.

Vuolteenaho, Tuomo, 2002, What Drives Firm-Level Stock Returns?, *Journal of Finance* 57, 233–264.

Whaley, Robert, 1986, Valuation of American Futures Options: Theory and Empirical Tests, *Journal of Finance* 41, 127–150.

Xing, Yuhang, Xiaoyan Zhang, and Rui Zhao, 2008, What Does Individual Option Volatility Smirk Tell Us about Future Equity Returns?, Working paper.

Table 1: Moneyness Category Definitions

Following Bollen and Whaley (2004), we define moneyness categories based on the delta of the stock, excluding options with deltas above 0.98 or below 0.02 in absolute value.

Category	Labels	Range
1	Deep in-the-money (DITM) call	0.875 $< \Delta_c <=$ 0.98
	Deep out-of-the-money (DOTM) put	-0.125 $< \Delta_p <=$ -0.02
2	In-the-money (ITM) call	0.625 $< \Delta_c <=$ 0.875
	Out-of-the-money (OTM) put	-0.375 $< \Delta_p <=$ -0.125
3	At-the-money (ATM) call	0.375 $< \Delta_c <=$ 0.625
	At-the-money (ATM) put	-0.625 $< \Delta_p <=$ -0.375
4	Out-of-the-money (OTM) call	0.125 $< \Delta_c <=$ 0.375
	In-the-money (ITM) put	-0.875 $< \Delta_p <=$ -0.625
5	Deep out-of-the-money (DOTM) call	0.02 $< \Delta_c <=$ 0.125
	Deep in-the-money (DITM) put	-0.98 $< \Delta_p <=$ -0.875

Table 2: Summary Table 1

This table presents some of the basic characteristics of the firms by predictability quintile during the 1970 − 2007 period. MSE rank represents the predictability quintile, with 1 being the most predictable (lowest mean-squared forecast error over the previous year) and 5 being the least predictable. Spread measures the bid-ask spread as a percentage of the closing stock price. Size is the market value of equity, and monthly turnover represents the average percentage of the shares outstanding that are traded each month.

MSE Rank		Spread	Size	Turnover	B / M	Leverage
1	Mean	1.7%	$ 5,547	6.8%	0.70	0.26
	Std	2.4%	$ 20,613	8.1%	0.50	0.19
	Max	82.6%	$ 446,879	249.3%	38.28	0.99
	Min	0.0%	$ 1	0.0%	0.00	-
	N	58106	76369	76363	75688	54998
2	Mean	1.8%	$ 3,998	8.3%	0.72	0.27
	Std	2.4%	$ 16,970	10.4%	0.56	0.21
	Max	102.7%	$ 518,079	372.8%	33.17	0.99
	Min	0.0%	$ 2	0.0%	0.00	-
	N	57928	75290	75290	74275	58043
3	Mean	1.9%	$ 3,364	9.9%	0.73	0.27
	Std	2.5%	$ 15,146	23.0%	0.65	0.22
	Max	76.0%	$ 578,128	3686.8%	33.46	0.99
	Min	0.0%	$ 1	0.0%	0.00	-
	N	57586	74259	74259	73043	59116
4	Mean	2.0%	$ 2,527	11.8%	0.72	0.25
	Std	2.8%	$ 11,531	15.1%	0.75	0.23
	Max	155.6%	$ 551,918	620.9%	46.67	0.99
	Min	0.0%	$ 1	0.0%	0.00	-
	N	57512	72743	72741	71076	60599
5	Mean	1.8%	$ 2,023	18.5%	0.60	0.20
	Std	2.7%	$ 8,942	21.8%	0.91	0.23
	Max	165.6%	$ 465,166	709.2%	43.72	1.00
	Min	0.0%	$ 0	0.0%	0.00	-
	N	58075	71596	71590	69499	63099

Table 3: Summary Table: Option Subsample

This table presents some of the basic characteristics of the firms by predictability quintile, which also have traded options in our sample, during the $1970 - 2007$ period. MSE rank represents the predictability quintile, with 0 being the most predictable (lowest mean-squared forecast error over the previous year) and 4 being the least predictable. Spread measures the bid-ask spread as a percentage of the closing stock price. Size is the market value of equity, and monthly turnover represents the average percentage of the shares outstanding that are traded each month.

MSE Rank		Spread	Size	Turnover	B / M	Leverage
1	Mean	0.7%	$ 12,601	11.4%	0.49	0.24
	Std	0.9%	$ 32,632	10.1%	0.34	0.17
	Max	11.4%	$ 446,879	246.0%	15.58	0.98
	Min	0.0%	$ 62	0.4%	0.00	-
	N	26105	26695	26695	26357	19959
2	Mean	0.8%	$ 8,940	14.0%	0.52	0.25
	Std	1.0%	$ 27,134	13.2%	0.42	0.19
	Max	12.2%	$ 518,079	372.8%	17.84	0.98
	Min	0.0%	$ 35	0.3%	0.00	-
	N	25525	26126	26126	25594	21393
3	Mean	0.8%	$ 7,516	16.5%	0.53	0.25
	Std	1.1%	$ 24,559	18.1%	0.51	0.21
	Max	18.0%	$ 578,128	1470.1%	29.52	0.98
	Min	0.0%	$ 29	0.2%	0.00	-
	N	25148	25676	25676	25098	21797
4	Mean	0.9%	$ 5,475	19.6%	0.55	0.23
	Std	1.3%	$ 18,524	18.4%	0.70	0.22
	Max	28.3%	$ 551,918	455.0%	46.67	0.98
	Min	0.0%	$ 6	0.3%	0.00	-
	N	24729	25234	25234	24448	22166
5	Mean	0.8%	$ 4,230	28.8%	0.44	0.17
	Std	1.0%	$ 13,951	25.5%	0.62	0.21
	Max	20.6%	$ 465,166	709.2%	17.79	1.00
	Min	0.0%	$ 14	0.2%	0.00	-
	N	26647	27125	27125	26195	24851

Table 4: Industries

Below we have the percentage of firm-months observations in our highest and lowest predictability quintiles for a sample of some the more extreme industries. Industries are designated according to Fama and French (1997).

	% of firm-months	
Industries	**Lowest MSE Quintile**	**Highest MSE Quintile**
Utilities	44.0%	5.3%
Books	31.9%	9.8%
Soda	31.3%	11.7%
Beer	27.8%	11.2%
Paper	24.1%	9.6%
Chips / Electronics	10.7%	31.3%
Computers	9.3%	29.1%
Personal Services	8.7%	32.2%
Health Care	8.7%	32.4%

Table 5: Forecast Mean Squared Error

Forecasting returns are based on a multivariate regression including past returns, profitability, leverage, and book to market. Firms are sorted into quintiles each month based on the mean-squared forecast error over the previous year. Returns are then forecasted for the subsequent 1,3,6, and 12 month horizon and the MSE of the forecasts are calculated for each quintile. This process is repeated monthly over the entire 1970 - 2006 sample period.

	Rank for MSE: 1 Month					Rank for MSE: 3 Month				
	1	2	3	4	5	1	2	3	4	5
Mean SE	0.0061	0.0096	0.0133	0.0196	0.0334	0.0168	0.0278	0.0403	0.0623	0.1242
Median SE	0.0017	0.0028	0.0039	0.0057	0.0093	0.0052	0.0094	0.0138	0.0207	0.0356
Std Error	0.0001	0.0001	0.0001	0.0002	0.0003	0.0002	0.0002	0.0003	0.0006	0.0017
	Rank for MSE: 6 Month					Rank for MSE: 12 Month				
	1	2	3	4	5	1	2	3	4	5
Mean SE	0.0292	0.0517	0.0773	0.1207	0.2979	0.0544	0.1023	0.1597	0.2581	0.9588
Median SE	0.0097	0.0185	0.0283	0.0425	0.0823	0.0193	0.0419	0.0679	0.1129	0.2423
Std Error	0.0002	0.0004	0.0006	0.0008	0.0035	0.0004	0.0007	0.0010	0.0015	0.0190

Table 6: Forecast MSE Two Way Sort

Similar to the previous table, however, now we sort by past forecast MSE and past one month variance. All within each row, recent volatility is held constant, and decreases in predictability are meet with statistically significant increases in MSE.

		1 Month Forecast					**6 Month Forecast**				
		Rank for Past 12 month MSE					Rank for Past 12 month MSE				
Mean SE		1	2	3	4	5	1	2	3	4	5
Std Error											
Previous Month Variance Rank	1	0.0026	0.0041	0.0056	0.0077	0.0127	0.0195	0.0350	0.0520	0.0782	0.1606
		0.0000	*0.0001*	*0.0001*	*0.0001*	*0.0002*	*0.0004*	*0.0007*	*0.0008*	*0.0013*	*0.0032*
	2	0.0037	0.0058	0.0079	0.0115	0.0198	0.0225	0.0410	0.0594	0.0920	0.2052
		0.0001	*0.0001*	*0.0001*	*0.0002*	*0.0003*	*0.0004*	*0.0007*	*0.0009*	*0.0014*	*0.0038*
	3	0.0047	0.0075	0.0104	0.0151	0.0259	0.0264	0.0459	0.0680	0.1084	0.2636
		0.0001	*0.0001*	*0.0001*	*0.0002*	*0.0004*	*0.0004*	*0.0007*	*0.0010*	*0.0016*	*0.0057*
	4	0.0063	0.0098	0.0136	0.0205	0.0353	0.0305	0.0535	0.0792	0.1293	0.3338
		0.0001	*0.0001*	*0.0002*	*0.0003*	*0.0006*	*0.0004*	*0.0008*	*0.0011*	*0.0021*	*0.0083*
	5	0.0123	0.0192	0.0273	0.0411	0.0723	0.0434	0.0751	0.1140	0.1721	0.5304
		0.0002	*0.0004*	*0.0005*	*0.0011*	*0.0017*	*0.0008*	*0.0012*	*0.0025*	*0.0035*	*0.0170*

Table 7: Forecast MSE Three Way Sort

Similar to the previous table, however, now we sort by past one year (excluding the past month) volatility, past month volatility, and finally on the forecast MSE ranking. Within each row, recent volatility is held constant, while past year volatility is held constant within each matrix. The first row is the average MSE and the second row, in italics, is the standard error.

		1 Month Forecast					**6 Month Forecast**				
		Rank for Past 12 month MSE					Rank for Past 12 month MSE				
Year Rank = 4		1	2	3	4	5	1	2	3	4	5
Previous Month Variance Rank	1	0.0109	0.0107	0.0117	0.0127	0.0154	0.0548	0.0704	0.0796	0.0817	0.0895
		0.0020	*0.0008*	*0.0004*	*0.0003*	*0.0008*	*0.0040*	*0.0046*	*0.0077*	*0.0033*	*0.0032*
	2	0.0091	0.0131	0.0143	0.0153	0.0185	0.0773	0.0819	0.0926	0.0946	0.1146
		0.0012	*0.0008*	*0.0006*	*0.0004*	*0.0008*	*0.0062*	*0.0038*	*0.0041*	*0.0035*	*0.0050*
	3	0.0098	0.0128	0.0147	0.0186	0.0219	0.1551	0.0948	0.1143	0.1140	0.1445
		0.0013	*0.0015*	*0.0005*	*0.0008*	*0.0009*	*0.0365*	*0.0056*	*0.0093*	*0.0039*	*0.0058*
	4	0.0116	0.0140	0.0188	0.0209	0.0253	0.1015	0.1318	0.1355	0.1477	0.1672
		0.0015	*0.0009*	*0.0008*	*0.0006*	*0.0010*	*0.0076*	*0.0089*	*0.0066*	*0.0054*	*0.0067*
	5	0.0140	0.0180	0.0204	0.0283	0.0316	0.2162	0.1761	0.2105	0.2064	0.2061
		0.0017	*0.0011*	*0.0009*	*0.0021*	*0.0013*	*0.0418*	*0.0144*	*0.0225*	*0.0104*	*0.0090*

Table 8: Panel Results

Regression output for 2 way fixed effects model run over weekly observations from February 1996 through December 2005. Rank is the MSE quintile used to measure predictability over 1 and 3 month horizons, where higher quintiles are associated with higher MSE and hence lower predictability. Return is the past 1 month return and return std is the realized standard deviation over the past month. Lag std is the realized standard deviation over the prior year excluding the past month. Industry std is the average of realized standard deviations for firms in that industry over the past month. Skewness and kurtosis are measured over the previous year based on daily closing prices. In order to match option maturity with return predictability horizon, observations for these regressions are limited to those options with less than 90 days until expiration.

	1 month	3 month	1 month	3 month	1 month	3 month	1 month	3 month
Return	(0.1120)	(0.0928)	(0.1261)	(0.1092)	(0.1232)	(0.1059)	(0.1235)	(0.1102)
	0.0060	0.0059	0.0059	0.0054	0.0058	0.0053	0.0059	0.0051
	(18.81)	**(15.79)**	**(21.49)**	**(20.24)**	**(21.06)**	**(19.96)**	**(21.09)**	**(21.69)**
Month	0.2394	0.2192	0.1891	0.1668	0.1883	0.1658	0.1931	0.1747
	0.0069	0.0067	0.0056	0.0053	0.0056	0.0053	0.0057	0.0050
	34.70	**32.82**	**33.73**	**31.69**	**33.46**	**31.43**	**33.88**	**35.07**
Rank	0.0127	0.0134	0.0049	0.0050	0.0050	0.0052	0.0043	0.0049
	0.0008	0.0008	0.0007	0.0006	0.0007	0.0006	0.0006	0.0007
	15.56	**17.40**	**7.45**	**8.31**	**7.67**	**8.57**	**6.77**	**6.88**
Leverage	0.1147	0.1240	0.0756	0.0841	0.0734	0.0819	0.0771	0.0846
	0.0166	0.0158	0.0127	0.0116	0.0127	0.0116	0.0123	0.0109
	6.91	**7.86**	**5.98**	**7.24**	**5.78**	**7.04**	**6.24**	**7.77**
Industry	0.1962	0.2092	0.1071	0.1108	0.1072	0.1107	0.0959	0.0978
	0.0191	0.0182	0.0137	0.0132	0.0136	0.0131	0.0132	0.0123
	10.29	**11.50**	**7.82**	**8.41**	**7.88**	**8.44**	**7.29**	**7.95**
Size	(0.0317)	(0.0376)	(0.0243)	(0.0292)	(0.0245)	(0.0293)	(0.0246)	(0.0291)
	0.0039	0.0041	0.0029	0.0030	0.0029	0.0030	0.0028	0.0028
	(8.14)	**(9.18)**	**(8.43)**	**(9.63)**	**(8.53)**	**(9.74)**	**(8.79)**	**(10.43)**
Year			0.4336	0.4519	0.4311	0.4496	0.4611	0.4794
			0.0138	0.0131	0.0139	0.0132	0.0140	0.0130
			31.43	**34.51**	**30.97**	**34.09**	**32.92**	**37.00**
Skew					(0.0030)	(0.0031)	(0.0033)	(0.0036)
					0.0009	0.0009	0.0007	0.0006
					(3.40)	**(3.66)**	**(4.74)**	**(5.64)**
Kurtosis							(0.0010)	(0.0012)
							0.0001	0.0001
							(11.85)	**(14.64)**
Fixed	Effects							
Time	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Table 9: Regression Results for Longer Horizons

Regression output for 2 way fixed effects model run over weekly observations from February 1996 through December 2005. Rank is the MSE quintile used to measure predictability over 1 and 3 month horizons, where higher quintiles are associated with higher MSE and hence lower predictability. Return is the past 1 month return and return std is the realized standard deviation over the past month. Lag std is the realized standard deviation over the prior year excluding the past month. Industry std is the average of realized standard deviations for firms in that industry over the past month. Skewness and kurtosis are measured over the previous year based on daily closing prices. In order to match option maturity with return predictability horizon, observations for these regressions are limited to those options with less than 90 days until expiration.

		Return	Month	Year	Rank	Leverage	Industry	Size	Skew	Kurtosis
6 month										
	Coefficient	(0.0977)	0.1578	0.4812	0.0050	0.0858	0.0829	(0.0309)	(0.0034)	(0.0012)
	SE	*0.0046*	*0.0049*	*0.0122*	*0.0006*	*0.0111*	*0.0120*	*0.0029*	*0.0006*	*0.0001*
	T stat	**(21.09)**	**32.43**	**39.30**	**8.65**	**7.74**	**6.93**	**(10.76)**	**(5.36)**	**(13.31)**
12 month										
	Coefficient	(0.0951)	0.1304	0.4548	0.0048	0.0817	0.0966	(0.0323)	(0.0034)	(0.0013)
	SE	*0.0072*	*0.0118*	*0.0139*	*0.0006*	*0.0113*	*0.0121*	*0.0029*	*0.0007*	*0.0001*
	T stat	**(13.13)**	**11.05**	**32.64**	**8.19**	**7.24**	**7.99**	**(10.95)**	**(4.60)**	**(14.34)**

Table 10: Quintile Dummy Regression

This table demonstrates the empirical price impact of a 0.6% increase in implied volatility for each quintile decrease in predictability. The prices are generated by the standard Black-Scholes formula on a hypothetical stock with a spot price of $40. It is assumed that the most predictable quintile has an implied volatility of 30% annually and the risk free rate is 5%.

		Intercept	Return	Month	Year	Leverage	Industry	Size	Skew	Kurtosis	D2	D3	D4	D5
1 month														
	Coefficient	0.5539	(0.1235)	0.1931	0.4610	0.0771	0.0958	(0.0247)	(0.0033)	(0.0010)	0.0034	0.0084	0.0128	0.0171
	SE	*0.0543*	*0.0058*	*0.0057*	*0.0140*	*0.0125*	*0.0132*	*0.0028*	*0.0007*	*0.0001*	*0.0014*	*0.0017*	*0.0021*	*0.0028*
	T stat	**10.20**	**(21.13)**	**33.85**	**32.85**	**6.19**	**7.28**	**(8.82)**	**(4.74)**	**(11.84)**	**2.40**	**4.77**	**5.94**	**6.14**
3 month														
	Coefficient	0.6143	(0.1110)	0.1685	0.4768	0.0831	0.0973	(0.0277)	(0.0036)	(0.0011)	0.0032	0.0079	0.0137	0.0172
	SE	*0.0543*	*0.0054*	*0.0053*	*0.0133*	*0.0113*	*0.0126*	*0.0028*	*0.0006*	*0.0001*	*0.0012*	*0.0015*	*0.0019*	*0.0027*
	T stat	**11.32**	**(20.57)**	**31.86**	**35.96**	**7.38**	**7.75**	**(9.77)**	**(5.54)**	**(13.38)**	**2.56**	**5.17**	**7.13**	**6.51**
6 month														
	Coefficient	0.5872	(0.1024)	0.1530	0.4682	0.0795	0.0812	(0.0256)	(0.0036)	(0.0011)	0.0018	0.0084	0.0133	0.0194
	SE	*0.0500*	*0.0043*	*0.0049*	*0.0122*	*0.0108*	*0.0114*	*0.0026*	*0.0006*	*0.0001*	*0.0011*	*0.0014*	*0.0017*	*0.0025*
	T stat	**11.75**	**(23.55)**	**31.47**	**38.43**	**7.37**	**7.12**	**(9.69)**	**(5.78)**	**(12.38)**	**1.57**	**6.01**	**7.68**	**7.77**
12 month														
	Coefficient	0.6094	(0.0922)	0.1345	0.4644	0.0772	0.0821	(0.0264)	(0.0030)	(0.0012)	0.0025	0.0076	0.0119	0.0182
	SE	*0.0487*	*0.0050*	*0.0050*	*0.0130*	*0.0109*	*0.0110*	*0.0026*	*0.0006*	*0.0001*	*0.0011*	*0.0014*	*0.0017*	*0.0024*
	T stat	**12.52**	**(18.25)**	**26.67**	**35.78**	**7.11**	**7.45**	**(10.14)**	**(5.19)**	**(15.08)**	**2.25**	**5.54**	**6.98**	**7.56**

Table 11: Price Impact

This table demonstrates the empirical price impact of a 0.6% increase in implied volatility for each quintile decrease in predictability. The prices are generated by the standard Black-Scholes formula on a hypothetical stock with a spot price of $40. It is assumed that the most predictable quintile has an implied volatility of 30% annually and the risk free rate is 5%.

	¡- Most Predictable to Least Prdictable -¿				
Strike Price	1	2	3	4	5
	Panel A: Time to maturity = 7 days				
30	$ 10.029	$ 10.029	$ 10.029	$ 10.029	$ 10.029
35	$ 5.034	$ 5.034	$ 5.034	$ 5.034	$ 5.034
40	$ 0.682	$ 0.693	$ 0.704	$ 0.715	$ 0.726
45	$ 0.001	$ 0.002	$ 0.002	$ 0.002	$ 0.002
50	$ 0.000	$ 0.000	$ 0.000	$ 0.000	$ 0.000
	Panel B: Time to maturity = 91 days				
30	$ 10.415	$ 10.420	$ 10.425	$ 10.430	$ 10.436
35	$ 5.920	$ 5.943	$ 5.966	$ 5.990	$ 6.014
40	$ 2.629	$ 2.669	$ 2.708	$ 2.747	$ 2.787
45	$ 0.896	$ 0.929	$ 0.962	$ 0.996	$ 1.029
50	$ 0.240	$ 0.257	$ 0.274	$ 0.292	$ 0.310
	Panel C: Time to maturity = 182 days				
30	$ 10.962	$ 10.979	$ 10.996	$ 11.013	$ 11.031
35	$ 6.889	$ 6.928	$ 6.968	$ 7.008	$ 7.048
40	$ 3.848	$ 3.903	$ 3.958	$ 4.013	$ 4.068
45	$ 1.925	$ 1.979	$ 2.032	$ 2.086	$ 2.140
50	$ 0.877	$ 0.917	$ 0.958	$ 0.999	$ 1.041
	Panel D: Time to maturity = 273 days				
30	$ 11.542	$ 11.569	$ 11.597	$ 11.625	$ 11.654
35	$ 7.737	$ 7.788	$ 7.839	$ 7.890	$ 7.942
40	$ 4.829	$ 4.895	$ 4.962	$ 5.028	$ 5.095
45	$ 2.831	$ 2.899	$ 2.967	$ 3.035	$ 3.103
50	$ 1.576	$ 1.634	$ 1.693	$ 1.753	$ 1.812
	Panel E: Time to maturity = 365 days				
30	$ 12.113	$ 12.149	$ 12.185	$ 12.222	$ 12.260
35	$ 8.499	$ 8.559	$ 8.619	$ 8.680	$ 8.741
40	$ 5.684	$ 5.759	$ 5.835	$ 5.911	$ 5.987
45	$ 3.650	$ 3.729	$ 3.809	$ 3.888	$ 3.968
50	$ 2.269	$ 2.342	$ 2.415	$ 2.489	$ 2.563

Table 12: Bollen Whaley Comparison

As a check against previous work, we replicate the results presented in table 8 of Bollen and Whaley (2004). We use the same sample of 20 firms, but are limited by data availability to the period 1996-2000 in contrast to 1995-2000 in their results. Given the broad nature of our our database (over 700 firms appear in our sample), daily observations for the life of an option are uncommon, so we rehedge weekly, rather than daily. The results shown in the top panel correspond to the mean monthly abnormal return, over the 20 large firms in Bollen and Whaley's sample, to selling an option and delta hedging with daily rebalancing it until expiration. The lower panel shows the same 20 firms with data from Optionmetrics and delta hedged weekly.

	Delta Value Categories				
	1	2	3	4	5
	Average Abnormal Returns				
BW Mean (1995-2000)	0.0140	0.0016	0.0019	0.0037	0.0183
Allred Mean (1996-2001)	0.0143	0.0031	0.0010	0.0061	0.0106

Table 13: Delta Hedged Returns

Below we have the abnormal returns, as defined in Bollen and Whaley (2004), to selling options by moneyness category, arranged from left to right, and predictability ranking, going from most predictable at the top to least at the bottom. Panel A contains the means abnormal return for the entire sample, panel B has means from the first half of the sample, panel C the medians from the entire sample, and finally panel d has the means from the second half of the sample. Options that are ITM or DITM are excluded from the results consistent with Bakshi and Kapadia (2003)

Panel A: Full Sample	Moneyness Category					Panel B: 1996-2001	Moneyness Category				
Mean	1	2	3	4	5	Mean	1	2	3	4	5
1	0.032	0.016	0.004	0.036	0.061	1	0.031	0.015	0.005	0.043	0.073
2	0.058	0.029	0.006	0.050	0.092	2	0.057	0.026	0.006	0.058	0.109
3	0.079	0.046	0.011	0.070	0.118	3	0.077	0.046	0.013	0.080	0.135
4	0.135	0.068	0.015	0.100	0.177	4	0.146	0.072	0.017	0.120	0.217
5	0.209	0.108	0.025	0.159	0.296	5	0.232	0.118	0.028	0.189	0.385

Panel C: Full Sample						Panel D: 2002-2007					
Median	1	2	3	4	5	Mean	1	2	3	4	5
1	0.030	0.019	0.014	0.037	0.055	1	0.033	0.017	0.002	0.029	0.051
2	0.053	0.033	0.017	0.053	0.084	2	0.059	0.032	0.006	0.043	0.075
3	0.083	0.049	0.023	0.068	0.109	3	0.081	0.046	0.009	0.058	0.102
4	0.132	0.070	0.026	0.094	0.160	4	0.124	0.064	0.012	0.076	0.139
5	0.210	0.107	0.036	0.142	0.266	5	0.171	0.092	0.020	0.108	0.166

Table 14: Delta Hedged at bid-ask

Below we have the abnormal returns, as defined in Bollen and Whaley (2004), to selling options by moneyness category, arranged from left to right, and predictability ranking, going from most predictable at the top to least at the bottom. Options that are ITM or DITM are excluded from the results consistent with Bakshi and Kapadia (2003). Options are sold during the fourth week of each month at the bid and repurchased at the ask. Delta hedging in the underlying is done such that the trader must buy at the ask and sell at the bid.

Panel A: Full Sample	Moneyness Category				
	1	2	3	4	5
1	0.020	0.001	-0.015	0.013	0.030
2	0.033	0.009	-0.018	0.022	0.055
3	0.049	0.020	-0.015	0.035	0.070
4	0.084	0.039	-0.013	0.058	0.101
5	0.154	0.076	-0.003	0.119	0.198

Table 15: Vega Hedged Returns

Below we have the abnormal returns, as defined in Bollen and Whaley (2004), to selling options by moneyness category, arranged from left to right, and predictability ranking, going from most predictable at the top to least at the bottom. Options that are ITM or DITM are excluded from the results consistent with Bakshi and Kapadia (2003). Options are sold during the fourth week of each month and immediately vega hedged using the option closest to ATM which has a maturity longer than the option being hedged. The net delta exposure is then hedged using the underlying. The values shown are means.

Panel A: Full Sample			Moneyness Category		
Means	1	2	3	4	5
1	0.027	-0.005	-0.015	-0.003	0.026
2	0.025	-0.018	-0.017	-0.005	0.043
3	0.020	-0.026	-0.021	-0.034	0.037
4	0.031	-0.052	-0.047	-0.062	0.025
5	-0.046	-0.104	-0.113	-0.100	0.046

Panel B: Full Sample					
Medians	1	2	3	4	5
1	0.003	0.007	0.001	0.019	0.027
2	0.016	0.007	0.004	0.026	0.056
3	0.019	0.004	0.005	0.027	0.054
4	0.068	-0.006	-0.005	0.026	0.084
5	0.070	-0.010	-0.011	0.020	0.139

Figure 1: Implied minus Realized

Similar to Figure 5 in Bollen and Whaley (2004), we calculate the difference between implied and realized volatility for all options between 1996 and 2001. Values below zero indicate that realized volatility is, on average, higher than implied. That fact that realized is higher than implied for close to the money options, yet delta hedged returns are often positive, is consistent with a negative volatility risk premium and ATM options having the highest vega.