Option Returns and the Cross-Sectional Predictability of Implied Volatility*

Alessio Saretto

The Krannert School of Management

Purdue University†

November 2005 ‡

Abstract

I study the cross-section of realized stock option returns and find an economically important source of predictability in the cross-sectional distribution of implied volatility. A zero-cost trading strategy that is long in straddles with a large positive forecast of the change in implied volatility and short in straddles with a large negative forecast produces an economically important and statistically significant average monthly return. The results are robust to different market conditions, to firm risk-characteristics, to various industry groupings, to options liquidity characteristics, and are not explained by linear factor models. Compared to the market prediction, the implied volatility estimate obtained from the cross-sectional forecasting model is a more precise and efficient estimate of future realized volatility.

*I thank Raffaella Giacomini, Richard Roll, Pedro Santa-Clara, and Walter Torous for their constant support, and Laura Frieder, Robert Geske, Amit Goyal, Mark Grimblatt and seminar partecipants at UCLA for valuable suggestions.

†West Lafayette, IN, 47907, phone: (765) 496-7591, e-mail: asaretto@purdue.edu.

‡The latest draft is available at: http://www.krannert.purdue.edu/faculty/saretto/.

1 Introduction

Volatility is central to the pricing of options as there is a one-to-one correspondence between the price of an option and the volatility of the underlying asset. In the context of Black and Scholes (1973), given the option price it is possible to obtain an estimate of the volatility by inverting the pricing formula. The resulting estimate is generally referred to as the implied volatility, which represents the market's estimate of the underlying future volatility over the life of the option. Options are essentially bets on volatility because an accurate prediction of future volatility delivers important economic information to traders. It is, therefore, not surprising that there is an extensive literature on predicting volatility. Granger and Poon (2003), for instance, survey the extant literature and broadly find that the market's forecast embedded in implied volatility is the best forecast of future volatility. However, this literature focuses mainly on predicting volatility of a single asset (frequently the S&P 500 index) using *time-series* methods. I show that there is important information in the *cross-section* of stock volatilities that leads to better forecasts of future volatility than those contained in the individual implied volatility itself. To the best of my knowledge, this is the first paper to study the predictability of the cross-section of individual equity option implied volatilities.

I use a system of Fama and MacBeth (1973) cross-sectional regressions for my forecasting exercise. Since the at-the-money implied volatility is directly linked to the underlying volatility it carries similar statistical properties. In particular it is very persistent. I estimate a mean-reversion cross-sectional model of implied volatility augmented with variables that improve the forecasting power (average of the historical realized volatility). I find that a stock with an at-the-money implied volatility below the cross-sectional average and below its own twelve months moving average, has a higher implied volatility in the next month. Similarly, a stock with an at-the-money implied volatility above the cross-sectional average and above its own twelve months moving average, has a lower implied volatility in the next month. Thus, cross-sectional regressions indicate a high degree of mean-reversion in implied volatilities.

I then study the economic implications of these forecasts through portfolio strategies. Specifically, I use the out-of-sample predictions produced by the cross-sectional forecasting model to construct a zero-cost trading strategy that involves a long position in a portfolio of options with large positive forecasts of the change in implied volatility and a short position in a portfolio of options with large negative forecasts. I study portfolios of calls, puts, and straddles and find that all these portfolios are quite profitable. After controlling for transaction costs and liquidity characteristics of the options the average monthly return to expiration of the straddles portfolio is as high as 8.8% with a monthly Sharpe ratio of 0.271. The returns of the underlying asset show no pattern across different portfolios, suggesting that the straddles portfolio returns are manifestations of superior volatility forecasts. The profitability of the long-short strategy is robust to different market conditions, to firm risk-characteristics, to various industry groupings, and to options liquidity characteristics. Moreover the strategy returns cannot be explained by CAPM or the Fama and French (1993) three-factor model augmented by the Carhart (1997) momentum factor or covariance with aggregate volatility.[1]

Since option prices are directly related to volatility forecasts, one can view the volatility forecasts as being alternative estimates of option prices. A volatility forecast that is higher than the implied volatility implies underpricing of the option, and a volatility forecast that is lower than the implied volatility implies overpricing of the option. It is interesting to explore the potential reasons for this success in identifying misspricing of options. One possible explanation for the profitability of the trading strategies is that the out-of-sample forecast produced by the model is a better estimate of future realized volatility than is the market forecast, represented by the implied volatility. I find evidence consistent with this hypothesis. Since the option's value is related to the underlying future volatility over the life of the option, this provides the underlying economic rationale for the trading profits reported earlier. As a further robustness check, I recompute the option prices by using my cross-sectional prediction of volatility rather than the market implied volatility. I then compute the return of the trading strategy using these modified prices. The profitability of the trading strategy completely disappears suggesting that the modified price obtained from the volatility forecast is closer to the 'true' price of the option. This provides further evidence of the economic importance of my results.

The market for equity options is active and has been constantly growing over the thirty years of its existence. The total volume of the equity options for the year 2004 was worth approximately 220 billion dollars. For comparison, the total volume of the

[1]Option payoffs are non-linearly related to payoffs of stocks. Therefore, a linear factor model is unlikely to characterize the cross-section of option returns. I use a linear model merely to illustrate that the option returns described in this paper are related to aggregate sources of risk in an obvious way.

S&P 500 index options was worth about 120 billion dollars.[2] There is also evidence that options traders are sophisticated investors: Easley, O'Hara, and Srinivas (1998) and Pan and Poteshman (2005) show that options' volume contains information about future stock prices. Given the non-negligible size of the market and the quality of option traders, it is useful to consider why it might be the case that option implied volatilities are predictable. One reason may be that the economic agents do not use all available information in forming expectations about future stock volatilities. In particular, they ignore the information contained in the cross-sectional distribution of implied volatilities and consider assets individually when forecasting their volatility. Stein (1955) shows that a combination of two different estimators produces a substantial gain in terms of the variance of the estimate. These shrinkage estimators are more efficient than the original estimates. The cross-sectional forecasts obtained in this paper have similar characteristics to those of shrinkage estimators, even though they are not formally constructed using Bayesian shrinkage techniques. This intuition provides the underlying statistical rationale for the trading profits reported earlier.

In a informationally efficient market, the implied volatility would be the best forecast of future volatility. Therefore, if the implied volatility is used to construct the forecast of future volatility, then the forecast should be no better than the implied volatility itself. To the contrary, the evidence presented in this paper suggests that the information contained in the forecast (based on readily available data) allows one to construct profitable trading strategies.[3] In the stock market context, reversal strategies (Jegadeesh (1990)) and momentum strategies (Jegadeesh and Titman (1993)) have been identified by Fama (1998) as those posing the most serious challenge to market efficiency. This paper extends this list by identifying the existence of similar strategies in the hitherto unexplored area of options markets.

The rest of the paper is organized as follows. The next section discusses the data. Section 3 contains a description of the implied volatility forecasting model and results of the estimation. In Section 4, I study the economic content of the volatility forecasts by analazying the performance of portfolios of options formed by sorting the forecasts. I examine the relation between the forecast and future realized volatility in Section 5.

[2]These figures are taken from the Options Clearing Corporation 2004 annual report, which can be found at http://www.optionsclearing.com/about/ann_rep/ann_rep_pdf/annual_rep_04.pdf

[3]Poteshman (2001) and Stein (1989) study the term structure of various estimates of implied volatility. Both these studies arrive at the conclusion that inefficiencies exist but do not explore their economic magnitude.

Section 6 concludes.

2 Data

The data on options are from the *OptionMetrics* Ivy DB database. The dataset contains information on the entire US equity option market and includes daily closing bid and ask quotes on American options as well as implied volatilities and deltas for the period between January 1996 and June 2005. (The implied volatilities and deltas are calculated using a binomial tree model *a la* Cox, Ross, and Rubinstein (1979).) Stock options are traded at the American Stock Exchange, the Boston Options Exchange, the Chicago Board Options Exchange, the International Securities Exchange, the Pacific Stock Exchange, and the Philadelphia Stock Exchange.

I apply a series of data filters to minimize the impact of recording errors. First I eliminate prices that violate arbitrage bounds. For calls, for example, I require that the option price does not fall outside the interval $(Se^{-\tau d} - Ke^{-\tau r}, Se^{-\tau d})$, where S is the value of the underlying asset, K is the option's strike price, d is the dividend yield, r is the risk free rate, and τ is the time to expiration. Second I eliminate all observations for which the ask is lower than the bid, or for which the bid is equal to zero, or for which the spread is lower than the minimum tick size (equal to \$0.05 for option trading below \$3 and \$0.10 in any other cases). More importantly, to mitigate the impact of stale quotes I eliminate from the sample all the observations for which both the bid and the ask are equal to the previuos day prices.

I construct time series of call, put and straddle returns for each stock in the sample. At any point in time equity options have traded maturities corresponding to the two near-term months plus two additional months from the January, February or March quarterly cycles. In order to have continuous time series with constant maturity I require the options to have expiration between 30 and 60 days (the second near-term month). Since all the options expire on the Saturday immediately following the third Friday of the expiration month, the criteria guarantees that all the option contracts selected have exactly the same maturity. Among those I then select the contracts which are closest to at-the-money (ATM). Since strike prices are spaced every \$2.5 apart when the strike price is between \$5 and \$25, \$5 apart when the strike price is between \$25 and \$200, and \$10 apart when the strike price is over \$200 it is not always possible to select option with exactly the desired moneyness. Options with moneyness lower than 0.95 or higher than 1.05 are therefore eliminated from the sample. The returns are constructed using as a reference price the average of the closing bid and ask quotes. Each month and for each stock, an option with the desired characteristics of moneyness and maturity is selected and the monthly return is computed using month beginning and month ending prices. At the beginning of the next month a new option with the same characteristics is selected and a new monthly return is calculated. Prices and returns for the underlying stock are taken from CRSP.

After applying the above mentioned filters, the sample is composed by 81,296 monthly observations. The average moneyness for calls and puts is very close to 1 while the average maturity is approximately 47 days. There are 4249 stocks in the sample for which is possible to construct at least one monthly return.

Summary statistics are reported in Table 1. Panel A shows that options returns are highly volatile and exhibit positive skewness and excess kurtosis. Calls have a positive mean return of 3.4%, respectively, while puts and straddles have a negative average returns of –3.6% and –0.3%. The median returns are instead all negative and large.

I report summary statistics for the ATM implied volatility (IV) and the annualized realized volatility (RV) of the underlying stocks in Panel B of Table 1. IV is computed as the average of the implied volatilities extracted from the call and the put contracts selected based on the maturity/moneyness filter. RV is computed as the standard deviation of daily realized returns from the beginning to the month to the expiration date (on average 47 days). IV and RV are close to each other and equal to 50.3% and 49.1% respectively. The overall distribution of RV is, however, more volatile and more positively skewed than that of IV. The average monthly change in both measures of volatility is very close to zero. Finally, IV surface exhibits a mild smirk — the 20% OTM put implied volatility (SmL) is 3.1% higher than the ATM volatility, while the 20% OTM call implied volatility (SmR) is 0.5% lower than the ATM volatility.[4]

Changes in implied volatility can be quite drastic and usually correspond to events of critical importance for the survival of a firm. For example the largest ΔIV in the sample is equal to 86% and corresponds to the release of particularly negative quarter loss for the fourth quarter of 1999 for UICI, a health insurance company. During the

[4]For a detailed discussion of the theoretical and empirical relation between the slope of the volatility surface and the properties of the risk-neutral distribution see Bakshi, Kapadia, and Madam (2003), Das and Sundaram (1999), Dennis and Mayhew (2002), and Toft and Prucyk (1997).

month of December, UICI options went from trading at an ATM implied volatility of 31% to an IV of 117%. The stock price lost 56% of its value in the same month. Many of the other large spikes in volatility happen during months of large declines in stock prices. For example, the implied volatility of the stocks in the technology sector jumped over 150% during the "burst" of the Nasdaq bubble in the spring of 2000. Considering stocks individually, spikes in implied volatility also happen in occasion of earnings announcements (Dubisnky and Johannes (2005)).

Individual equity options share some characteristics with index options, which have been the primary subject of research. Figure 1 plots the time series of VIX (volatility of S&P 500 index) and the time series of the cross-sectional average implied volatility (IMP). Naturally the level of IMP is much higher than that of VIX. Both series have spikes that correspond to important events, such as the Russia crisis of September 1998. The two variables are also highly correlated. The correlation coefficient of the changes in VIX and IMP is 53% when IMP is computed using an equally-weighted average and 73% when a value-weighted average is instead used.

However, the two variables differ in an important way – The average stock implied volatility is more persistent. The autocorrelation coefficient of the average implied volatility is equal to 0.964; the same coefficient is 0.755 for VIX. This high degree of persistence in stock implied volatility is the central feature of the forecasting model for stock's implied volatility that is developed later in the paper (see Section 3). Another way in which the equity option market differs from the index option market is that the asymmetric volatility effect of Black (1976) is less pronounced. The monthly correlation between the underlying asset return and change in implied volatility is –0.52 for index options and –0.32, on average, for individual stocks (see for example Dennis, Mayhew, and Stivers (2005)).

3 Cross-Sectional Predictability of Volatility

Since options are essentially bets on volatility, a predictive model of the future implied volatility will have economic significance. In this section, therefore, I start by constructing a forecasting model for the options implied volatility.

3.1 Forecasting Model

The bulk of the finance literature has focused on predicting the volatility of a single asset using a time-series approach. I, on the other hand, use a cross-sectional regression to forecast the distribution of future implied volatilities. This approach is similar to Jegadeesh (1990) approach for identifying predictable patterns in the cross-section of stock returns. I specify the forecasting model as follows:

$$\begin{aligned} \Delta iv_{i,t+1} &= \alpha_t + \beta_t iv_{i,t} + \gamma_t X_{i,t} + \epsilon_{t+1} \\ X_{i,t} &= \{iv_{i,t} - \overline{iv}_{i,t-12:t-1}, iv_{i,t} - \overline{rv}_{i,t-12:t-1}\} \end{aligned} \tag{1}$$

where $iv_{i,t}$ is the natural logarithm of the ATM implied volatility for stock i measured at the beginning of month t, $\overline{iv}_{i,t-12:t-1}$ is the natural logarithm of the twelve months moving average of IV_i, and $\overline{rv}_{i,t-12:t-1}$ is the natural logarithm of the twelve months moving average of the realized volatility for stock i. The cross-sectional model avoids the problem of determining the unconditional mean of each stock's volatility, leaving a much simpler task of estimating the conditional cross-sectional average.

The model is motivated primarily by the existing empirical evidence, both at the aggregate and individual stock level, of a high degree of mean-reversion in volatility (see Granger and Poon (2003) and Andersen, Bollerslev, Christoffersen, and Diebold (2005) for comprehensive reviews). The twelve month average of implied volatility is included because lags below the first could contain some valuable information. I include an average rather than all the lags because the average is a conservative way of using all available information: it involves estimating only one parameter, instead of the twelve parameters corresponding to twelve lags and it does not lead to data loss when one of the lags is missing. Finally, realized volatility is included as an additional regressor because it could provide incremental information over implied volatility.

I choose to forecast the change in log volatility because that is the relevant variable in constructing option strategies. This is though equivalent to forecast the log of future volatility. To obtain the percentage change, in that case, it is necessary to subtract the log of the actual value of implied volatility from the regression fitted value. An alternative procedure would be to forecast the change in the level of future implied volatility. I chose to forecast the change in log volatility because in this way I avoid the problem of having to truncate the negative fitted values in the prediction of the level of volatility.

I proceed by estimating a Fama and MacBeth (1973) cross-sectional regression at each date t (the beginning of the month). I tabulate averages of the cross-sectional estimates and t-statistics adjusted for serial correlation in Table 2. I also report the average adjusted R^2 from the cross-sectional regression as a measure of in-sample performance and the moments of the cross-sectional distribution of Root Mean Square Errors (RMSE) as a measure of out-of-sample performance. I compute the root mean squared error as $RMSE(i) = \left[\frac{1}{T}\sum_{t=1}^{T}(\Delta iv^*_{i,t} - \Delta iv_{i,t})^2\right]^{0.5}$, where $\Delta iv^*_{i,t}$ is the predicted value based on the regression at time $t-1$ and $\Delta iv_{i,t}$ is the actual logarithmic change in implied volatility. RMSE measures, for each stock i, the deviation of forecasts from actual changes in volatility.

3.2 Cross-Sectional Regressions

Specification (1) in Table 2 shows that the future increase in volatility is negatively related to the current level of volatility, which implies that stocks' implied volatilities reverts towards the cross-sectional mean: a high level of volatility today, compared to the cross-sectional average, predicts a lower implied volatility in the future, or a negative change. The estimated coefficient of –0.069 is highly statistically significant and implies a cross-auto-correlation coefficient for the level of IV of 0.921 which is very close to the estimate obtained from the time series of cross-sectional averages.

Specification (2) shows that the coefficient on the difference between the current level of implied volatility and the twelve-month average is negative and largely significant, suggesting that the mean reversion property of the individual time-series of implied volatility is also an important factor in predicting the future change in IV. The richer dynamic of model (2) leads to an higher average adj-R^2 and a lower average RMSE. A similar conclusion can be drawn from specification (3) wherein the difference between the current level of implied volatility and the twelve-month moving average of realized volatility is considered. The best forecasting model among the four considered is reported in the fourth model wherein the two measures of past volatility are concurrently used as a predictor. In this model the change in implied volatility is negatively related to the last period IV and negatively related to both moving averages. The average adj-R^2 is quite large at 17.1%, and at times it is as high as 50%. Model (4) has also the lowest average

RMSE, equal to 14.2%, confirming that the better in-sample performance, measured by the adj-R^2, is not due to data overfitting. The forecasts produced by this model are used for the rest of the analysis in the paper.

One interesting result from the estimation of the forecasting model is that the accuracy of the predictions, either in and out of sample, is negatively related to changes in the aggregate level of volatility, as measured by VIX. The monthly correlation between the time series of adj-R^2 and changes in VIX is in fact equal to –0.38, while the correlation between the time series of cross-sectional averages of square deviations of the out of sample forecasts from the actual changes and changes in VIX is about 0.32. The movements of the cross-section of implied volatilities are more difficult to forecast when the aggregate volatility is increasing.

3.3 Portfolios Based on Cross-Sectional Forecasts

The regression model of the previous subsection indicates a high degree of predictability in the cross-sectional distribution of changes in implied volatility. However, the RMSE of the regression suggests that forecasts for individual stock volatility are still fairly noisy. One way of using the information from the cross-sectional regressions while maximizing the signal-to-noise ratio is to form portfolios. Specifically, for each month in the sample I form decile portfolios by ranking the out-of-sample volatility forecasts. Portfolio 10 is predicted to have the highest (positive) percentage change in implied volatility while portfolio 1 predicted to have the lowest (negative) percentage change in implied volatility. The advantage of forming portfolios is that a cardinal signal on predicted value is transformed into a more precise ordinal signal on the ranking of predicted values.

I report the results of this analysis in Table 3. I the first panel, for each portfolio, I report the average difference between IV and the cross-sectional average, $\hat{IV}$, and the difference between the IV and the twelve month moving average, $\overline{IV}$. The analysis confirm that the predictability in implied volatility is related to both the cross-sectional and the time-series mean reversion. For example, the average $IV - \overline{IV}$ for portfolio 1 is 15.0% and –11.8% for portfolio 10 – both are highly statistically significant, while the average $IV - \hat{IV}$ is 11.2% and –8.0%, respectively. The further the market estimate of future volatility is from the past measures, the more accurate is the prediction for the change in future volatility. I also report the mean of the out-of-sample predicted change (in levels), ΔIV^* as well as of the actual change, ΔIV. The two averages are remarkably close. For example, the average forecast for portfolio 1 is –7.3% while the actual change is –5.9%. Similarly the forecast for portfolio 10 is 3.9% while the actual change is 3.5%.

4 Economic Significance

Since changes in volatility are the most important determinant of option returns, the predictability uncovered in the previous section can potentially be economically significant to investors. In this section I investigate this possibility by analyzing the performance of portfolios of options formed on the basis of the implied volatility forecasts.

I compute a prediction of each stock's future percentage change in volatility at the beginning of each month in a real time fashion. I use implied and realized volatility measures available at the beginning of the month and parameter estimates obtained from the previous month to obtain predictions of expected changes of IV. I then form decile portfolios by ranking these out-of-sample volatility forecasts. Portfolio 10 is predicted to have the highest (positive) percentage change in implied volatility while portfolio 1 is predicted to have the lowest (negative) percentage change in implied volatility. Equally-weighted monthly returns on calls, puts, straddles, and underlying stocks of each portfolio are computed and the procedure is then repeated for every month in the sample. On average the portfolios are composed of 134 different equities in each month.

Table 4 reports the results of this exercise. The pattern in the portfolio average returns is in line with the predicted change in volatility. Portfolio 1, which corresponds to the highest predicted decrease in volatility, has negative average returns equal to –3.1%, –12.0%, and –7.2% for calls, puts, and straddles, respectively. The average returns increase monotonically as one goes from portfolio 1 to 10. Portfolio 10, which corresponds to stocks with large predicted increase in volatility, has average returns equal to 8.8%, 5.6%, and 5.6% for calls, puts, and straddles, respectively. The call and put portfolios are, however, characterized by very high volatility that ranges from 41% to 53% per month. The volatility of the straddle portfolios is instead much lower and between 12.3% and 17.5% per month.

I also report two measures related to the risk-return trade off for the portfolios: Sharpe ratio (SR) and certainty equivalent (CE). CE is computed for a long position in the portfolio and is constructed using a power utility with coefficient of relative risk aversion equal to 3.[5] For reference, the market, as proxied by the value-weighted CRSP portfolio, has a Sharpe ratio of 0.103 and a CE of 0.45% per month. Because of the high volatility and extreme minimum and maximum returns, which imply large high order moments, all call and put portfolios have low SR and negative CE. Straddle portfolios, on the other hand, have high Sharpe ratios and certainty equivalents.

The returns to a long-short strategy, that is long in portfolio 10 and short in portfolio 1, are noteworthy. The call and put strategies have high average returns and volatility that is generally lower than the two originating portfolios leading to large monthly Sharpe equal to 0.334 and 0.683 for calls and puts respectively. The very large minimum returns, –107% however, lead to very low CE for the call long-short portfolio. The put strategy has still a large negative minimum return, –50.3%, but a positive CE. In contrast the long-short staddle strategy has an average return of 12.9% with a 13.2% monthly standard deviation (the minimum monthly return in the sample is –21.6%). This leads to a monthly Sharpe ratio of 0.977 and a certainty equivalent of 10.5% per month.

Please note that these option returns do not appear to be driven by directional exposure to the underlying asset. When underlying stocks are sorted according to the same portfolio classification, the returns of the stock portfolios do not exhibit any particular trend in either the first or the second moment. Altogether the evidence confirms that the ability to forecast volatility leads to predictability in option returns. The cross-section of option returns appears to be very well explained by the cross-section of volatility forecasts. The long-short portfolio option return is statistically significant and economically large.

Note also that the results do not appear to be driven by microstructure effects. Since the returns are computed from the mid-point prices they are not affected by the bid-ask bounce effect of Roll (1984). As an empirical validation I repeat the analysis using option returns computed from the second, as opposed to the first, to the last trading day of the month. In this case to eliminate stale quotes, I do not consider bid and ask prices that do not change between the first and the second trading day of the month. Untabulated results indicate that a very small fraction of the returns reported in Table 4 are due to the first day of trading. For example, the effect of skipping one day in the return construction alters the straddles long-short strategy average return by 1.3%, from 12.9% to 11.6% respectively. As the second return corresponds to a holding period of 21

[5] CE is potentially a better measure of the risk-return trade off then SR because it takes into account all the moments of the return distribution.

instead of 22 days, everything else the same, the expected difference in the return should be 0.6%. The unaccounted difference, 0.7%, is then due to the process of lagging by one day the portfolio construction from the calculation of the signal (expected change in volatility).

4.1 Robustness

I implement a series of robustness checks that are designed to help us understand the profitability of these portfolios. The previous section shows that the straddle returns are most clearly related to the source of predictability, that is changes in future volatility. For this reason, I focus only on the straddle portfolios.[6]

4.1.1 Sub-Samples

I replicate the analysis of Table 4 by dividing the data into two sub-samples. The sub-samples are formed by considering different states of the market and of the aggregate volatility. The states are determined by the sign of the market value-weighted CRSP portfolio returns and by the sign of the changes in the VIX index. Mean returns and t-statistics are reported in Table 5. In Panel A the sample is divided according to the size of the market return. The return difference between positive and negative market periods is decreasing with the size of the predicted change in implied volatility. Portfolio (1) performs better when the market is rising, while the opposite is true for portfolio (10). For the long-short portfolio the spread between up and down market is equal to –4.8% and marginally statistically significant.

Panel B of Table 5 shows how the portfolio returns differ in periods of increasing and decreasing aggregate volatility (VIX). The conditional portfolio returns are higher in months in which VIX is increasing. The return difference increases moving from portfolio (1) to portfolio (10). The average return difference is virtually zero for portfolio (1) and 9.2% for portfolio (10). In seven of the ten cases this difference is significantly grater than zero. The long-short strategy has returns of 17.6% in months of positive changes in VIX and 8.5% in months of negative changes in VIX. The difference is economically important at 9.2% per month (and statistically significant). Straddles have a positive

[6]Complete results for portfolios of calls and puts, which are qualitatively similar, can be obtained from the author upon request.

return when volatility increases, so it is natural that the portfolios have a greater return when aggregate volatility is moving up. What the result says however is that this seems to be particularly the case for options with a positive forecasts or, in other terms, a lower implied volatility.

Untabulated results for portfolios of calls and puts differ in only one dimension: the call long-short portfolio has a much larger average return (19.5%) when the market is up versus 10.7% when the market is down. The put long-short strategy is more profitable (18.1%) when the market is down versus 9.7% when the market is up. This result, is however, expected because the option prices move in accordance with the direction of underlying asset. When aggregate volatility is considered as the reference state variable I observe the same patterns: call portfolios are all negative when volatility is rising while put portfolios are all negative when volatility is decreasing. This of course is due to the negative correlation between the aggregate volatility and the stock market. When volatility goes up the market is tanking which makes call returns negative and put returns positive.

4.1.2 Stock Characteristics

Since options are derivative securities, it is reasonable to assume that option returns depend on the same sources of risk that characterize stock returns. The absence of a formal theoretical model for the cross-section of option returns further warrants considering stock-risk factor related explanations for option returns. I, therefore, investigate how the long-short straddle returns are related to equity risk factors. I consider two-way independent sorts - one based on volatility forecast and the second based on firm characteristics. The characteristics chosen are beta, size, book-to-market and past return. The first three of these are motivated by Fama and French (1992) while the last one is due to evidence of momentum profits by Jegadeesh and Titman (1993).[7] I sort stocks into quintile portfolios, as opposed to decile, to keep the portfolios well populated. Breakpoints for size, book-to-market, and momentum are obtained from Kenneth French's website.

Table 6 shows the results of these double sorts. The profits due to volatility predictability persists in any beta, size, book-to-market, and momentum portfolios indicat-

[7]See also Amin, Coval, and Seyhun (2004), who find a relation between index option prices and momentum.

ing that the "volatility effect" is not subsumed by other effects typical of the cross-section of stock returns. The long-short portfolio strategy has statistically significant average returns that range from 6.9% to 12.8% per month — Note that these are quintile portfolios. Moreover the magnitude of the long-short portfolio returns seems to be related to the firm characteristics: the average returns are higher for small market capitalization, high book-to-market ratio and past loser stocks. Note that all these characteristics are also related to volatility: anything else constant stocks with small size, high book-to-market ratio, and poor past market performance tend to be high volatility stocks.

At the same time, the empirical regularities found in stock returns also migrate to option returns. Stocks with high beta, small size, high book-to-market, and poor stock market performance all have larger average straddle returns. However the magnitude of the difference between the top and bottom portfolios is not always statistically significant, and it is not large enough to explain the "volatility predictability" phenomenon. Nonetheless it appears that the characteristics positively covary with the size of the volatility forecasts. The return difference between portfolios of options with high and low stock beta, small and large market capitalization, high and low equity book to market, and negative and positive past stock performance tends to be positively related to the implied volatility forecast. For example, the S–B size portfolio has an average return of 7.6% per month in the top quintile of volatility forecasts and 3.9% per month in the bottom quintile of Δiv^*.

Figure 1 shows that the equity option market was particularly active during the years of the "technology bubble." It is, therefore, imperative to establish if the volatility predictability is a phenomenon in only the technology industry. In unreported results, I find this not to be the case. The long-short volatility strategy is quite profitable in each industry. The highest average return (9.4% per month) is in the finance sector while the lowest return (6.9%) is in the utilities industry.

I conclude that the option returns do covary with the same stock characteristics that are found to be important for stock returns, but this covariance is not enough to explain the portfolio returns based on the volatility predictability.

4.2 Risk Adjusted Returns

I proceed by examining whether the profitability of the volatility option portfolios is related to aggregate risk. I regress the straddle portfolio returns on various specifications of a linear pricing model composed by the Fama and French (1993) three factors model, the Carhart (1997) momentum factor, and changes in the aggregate implied volatility, as measured by VIX. Estimated parameters for portfolio (1) (lowest predicted increase in volatility), (10) (highest predicted increase in volatility) and (10-1) are reported in Table 7. Of the five factors the momentum factor does not seem to play a very important role. The loadings of the remaining factors, with the exception of the volatility changes, are higher for portfolio (1) than they are for portfolio (10). There is however a lot of noise in these estimates and the t-statistics tend to be quite low. On the contrary, the linear factor specification does a much better job at explaining the returns of the long-short portfolio. With the exception of momentum, the other factors have significant loadings and the adjusted-R^2 are quite large, especially considering that we are analyzing a portfolio of options.

Note however that from the qualitative standpoint the results appear quite surprising. The signs of the estimated parameters do not suggest that the option returns, and in particular the returns of the long-short strategy, are explainable in terms of remuneration for risk. On the contrary, the return of portfolio (10-1) are negatively related to movements in the three stock market factors and are positively related to the changes in aggregate volatility. The long-short strategy appears therefore to be quite attractive, even without considering the extremely large returns, because it also hedges the sources of aggregate risk that are priced in the stock market. This is also captured by the fact that the risk-adjusted return, $\hat{\alpha}$, is higher than the strategy average return.

4.3 Multiple Maturities and Horizons

The analysis conducted so far is based on the one-month prediction of the future change in volatility and on the return of options with less than two month to maturity. On one hand options on the same underlying but with different maturity do not always share the same characteristics. On the other hand, the evidence that implied volatility follows a persistent process (estimates of the half-life of volatility in general lay around the 6 months time period) suggests that the longer horizon predictability might be stronger than that at one-month horizon. In this section, therefore, I investigate to what extent the maturity of the contracts and the horizon of the volatility forecasts impacts the profitability of the long-short trading strategy. I consider three different maturities and horizons: 1, 3 and 6 months.[8] Summary statistics for the long-short straddle portfolios are reported in Table 8. In the first panel I report summary statistics for straddles portfolios with different maturities obtained by sorting the 1 month volatility forecasts (same as Section 3). The three columns differ in the choice of the options used to form the portfolios: in the first column (*1-month*) I only consider options that expire between 1 and 2 months (same as before), in the second options that expire between 3 and 4 months, and in the last options that expire between 6 and 7 months. All the returns are then computed over the next month. As it was expected the table shows that prices of longer-maturity options are less sensible to short-run changes in volatility. Consequently the returns are smaller when longer maturity options are considered — Note however that the Sharpe ratios do not decrease as much as the average returns, indicating that a long-short portfolio of longer maturity options might not be always ignored, especially when the volatility of the strategy is also of interest.

In the second panel, I instead report summary statistics for three long-short portfolios constructed by ranking volatility forecasts over the different horizons and by computing returns as buy and hold over the particular horizon. So for example the 6 months long-short strategy is obtained by shorting 6 months ahead out-of-sample volatility forecasts. Once it is determined which options go in the long and short portfolio, the options are bought and sold and kept in the portfolio for 6 months. The procedure is repeated each month, so that at any point in time there exist six long and six short portfolios. Each of these six portfolios contains options which will mature at a different point in time. The returns are realized at end of the month before expiration. Therefore the six-month horizon portfolio has monthly return observations, but each is a six months holding period return.

The first step in the procedure involves re-estimating the forecasting model (1) by considering three and six months horizons change in volatility, defined as $\Delta iv_{t+3} = \log(IV_{t+3}) - \log(IV_t)$ and $\Delta iv_{t+6} = \log(IV_{t+6}) - \log(IV_t)$. The estimates of these regressions are not reported. As expected, the coefficient on IV_{t-1} becomes more negative with

[8] For convenience and to keep the notation of the table as simple as possible, I refer to 1, 3, and 6 months to maturity. These maturities are in reality more than what they are noted, because the contracts expire on the third Friday of the following month, so that the proper notation would have to be less than 2, 4, and 7 months.

horizon, while the coefficients on moving average of past volatility become larger. The average adj-R^2 increases with the horizon, from 12.9% to 14.9% and 15.0%, respectively. The average holding returns however do not increase with horizon.

This exercise serves as a reality check to understand whether the results of previous sections are entirely driven by a short-run phenomenon, and hence controlling that the volatility predictability extends to the entire term-structure of implied volatilities. Combining the results of the two panels I conclude that the larger part of the returns is realized in the first month and hence that the implied volatility predictability has economic implications in the very short-term.

4.4 Trading Execution

There is a large body of literature that documents the finding that transaction costs in the options market are quite large and are in part responsible for some pricing anomalies, such as violations of the put-call parity relation.[9] For this reason, it is essential to understand to what degree these frictions prevent investor from exploiting the volatility predictability uncovered in this paper. Therefore, in this section I discuss the impact of transaction costs, measured by the bid-ask spread, and margin requirements on the feasibility of the long-short strategy.

I consider the costs associated with executing the trades at prices inside the bid-ask spread. The results reported so far are based on returns computed using the mid-point price as a reference; however it might not be possible to trade at that price in every circumstance. De Fontnouvelle, Fisher, and Harris (2003) and Mayhew (2002) document that the effective spreads for equity options are large in absolute terms but smaller relative to the quoted spreads. Typically in fact the ratio of effective to quoted spread is less than 0.5. Moreover Chan, Chung, and Johnson (1995) study the intraday behavior of option prices and document that spreads tend to increase at the end of the day. Since transactions data is not available for the sample considered in this study, I consider an effective spread equal to 75% of the quoted spread. Given the evidence cited above, this transaction cost estimate should be a conservative measure of the real execution costs.

[9] See for example Gould and Galai (1974), Figlewski (1989), Ho (1984), George and Longstaff (1993), Ofek, Richardson, and Whitelaw (2004), Santa-Clara and Saretto (2005), and Swidler and Diltz (1992).

Since transaction costs are very large in the option market, a real-world investor would not choose to reverse his positions at the end of the month, but would wait to unwind the strategy at the expiration date of the options. To account for this I compute the return to expiration, which corresponds to an average holding period of 47 days. The average return to expiration of the long-short strategy is equal to 16.1% when the returns are computed from mid-point prices and 5.5% when transaction costs are applied, suggesting that the volatility predictability originates some profitable opportunities. Moreover, to eliminate the concern that the results might come from options that are extremely thinly traded I repeat the analysis by splitting the sample in different liquidity groups which are obtained by ranking stocks on the base of the liquidity characteristics of the options. For each stock I compute the average quoted bid-ask spread of all the options series traded in the previous month as well as the daily average dollar volume. I then sort stocks based on these characteristics and report average returns and t-statistics for the long-short portfolio. In the top panel of Table 9 I tabulate results for terciles (low, medium and high) obtained by sorting the average bid-ask spread; in the bottom panel I separate stocks based on the options daily average dollar volume. I report the average return computed from the mid-point price (MidP) and from the bid and ask prices (EBA) adjusted for an effective bid-ask spread eqaul to 75% of the quoted spread.

Altogether the results indicate that the large return of the long-short strategy does not stem from very illiquid options: while it seems that options with large quoted spreads tend to have a higher returns than options with small spreads (16.6% versus 13.3%) it is not clear that the same result is supported by the volume portfolios. Moreover, in both cases the highest average return is realized in the medium tercile portfolio. Taking trasnaction costs into account the portfolio returns decrease substantially but remain sizeble. For example in the case of the medium bid-ask spread portfolio the average return is 8.6% after trading costs. Lagging the portfolio formation by one day, as discussed in section 4, has a small impact on the portfolio returns suggesting that the results are not due to micro-structure effects.

Santa-Clara and Saretto (2005) show that margin requirements on short-sale positions can be quite effective at preventing investors to take advantage of large profit opportunities in the S&P 500 option market. However margins on short positions have a smaller impact on trades that involve options with strike prices close to the money, as it is the case of the strategy outlined in the previous section. The short side of the long-short strategy involves options with high IV and low (negative) expected percentage change in IV. Therefore, these options have high prices and relatively high price-to-underlying ratios. Margin requirements for these options are relatively low and should not affect the execution of the strategy.

5 Cross-Sectional Forecast as a Shrinkage Estimator

If economic agents consider assets individually when forecasting their volatilities they will not produce the most efficient predictions. It is in fact a principle of Bayesian statistics that a way to improve the accuracy, defined as the expected square error, of an estimate is to combine that estimate with some informative prior. The way the two estimates, the original and the prior, are then combined is by shrinking towards the prior the estimates that are very far from it. This type of estimators are referred to as *Shrinkage estimators*.[10]

A possible explanation to the profitability of the trading strategies is that the forecast produced by the cross-sectional model is a better estimate of future realized volatility than the market forecast, represented by the implied volatility. The cross-sectional estimate, CS from now on, is computed by adding up the out-of-sample prediction of the change in the level of IV, $\Delta IV_t^* = IV_t \times \Delta iv_t^*$, to the current level of implied volatility. In other words CS is given by the combination of two different forecasts: one produced by the market and another computed using the cross-sectional forecasting model 1

$$CS_t = IV_t + \underbrace{\Delta IV_t^*}_{shrinkage factor}$$

Despite the way it is obtained, which is very different from the Bayesian approach, ΔIV_t^* shares some property with the shrinkage factor of the Stein's estimator. Indeed it operates like a shrinkage factor in the sense that, as is evidenced in Table 3, it is negative for high values of IV and it is positive for low values of IV. In the ten portfolios

[10] Shrinkage estimators were introduced by Stein (1955) and rely on the idea that a combination of two different estimators might produce a substantial gain in terms of the variance of the estimate. Efron and Morris (1977) present a non-technical discussion of the Stein's estimator while Morris (1983) offers a more rigorous presentation. For an example of application to the estimation of the variance covariance matrix see Ledoit and Wolf (2003), while probably the most common finance application of the shrinkage is the estimation of asset's betas, see for example Karolyi (1992).

constructed by ranking the out-of-sample log change in implied volatility, the average adjustment factor , ΔIV^*, is equal to –7.3% for the portfolio with the lowest predicted change, and it is equal to 3.9% for the portfolio with the highest predicted change. Note that respectively these portfolios have high, 64.8%, and low, 35.3%, current levels of implied volatility. ΔIV^* is therefore effectively shrinking the tails of the distribution towards the cross-sectional average. While in the classical Stein's estimator this property would be produced by explicitely operating on the cross-sectional average of implied volatility, the informative prior, here we obtain it by filtering the data through the estimated parameters of a cross-sectional regression. The analogy holds because of the property of the liner regression. The estimated parameters are such that the regression line fits through the cross-sectional average of the dependent variable, again the average of the stocks' IVs.

To offer more evidence in Panel A of Figure 2 I plot ΔIV^* versus the implied volatility as of the beginning of September 2001.[11] Similarly to the table, the figure shows that ΔIV^* tends to be negative for large values of the implied volatility, while the reverse is true for small values of IV. As a result the cross-sectional distribution of CS is less dispersed than that of IV in most of the months in the sample. In Panel B of Figure 2 I plot the tails of the non-parametric kernel density of the cross-sectional distribution of IV and CS. The left and right tails are plotted on the left and right hand-side, respectively. The "shrinkage" of the distribution of IV is quite apparent: the left tail of CS lays in fact at the right of the left tail of IV, while the opposite pattern is shown in the figure on the right. This is not an exception in the sample. Indeed in about 82% of the months the left tail of the cross-sectional distribution of CS lays to the right of the corresponding distribution of IV. Similarly in 92% of the months the right tail of CS is instead to the left of the right tail of IV. The range of CS, defined as the interval between the min and the max, lays within the range of the distribution of IV in 85% of the months. Moreover, the same range is on average 8.8% smaller than the range of the implied volatility distribution.

Having shown that CS has similar characteristics to the shrinkage estimator provides the underlying statistical rationale for the trading profits reported earlier. In the rest of the section I therefore rationalize the economic aspect of the reported predictability by investigating the relation between the two volatility forecasts, IV and CS, and the future

[11]The choice of the date is completely casual and corresponds to when I started the PhD program at UCLA.

realized volatility. In particular, using an approach similar to the study of Christensen and Prabhala (1998),[12] I test the hypothesis that CS is a better predictor, than the market forecast IV, of the future realized volatility. I run a horse race between CS and IV by estimating the following regression model

$$FV_{t,t+k} = \alpha + \beta X_t + \epsilon_t \tag{2}$$

where FV is the future realized volatility over the life of the option, and X is either the cross-sectional forecast, the implied volatility, or the past realized volatility, which is included as a benchmark case. The dependent variable is the realized volatility of daily returns over the life of the option. So, for example, if the implied volatility is extracted from a couple of options with 47 days to expiration I compute the annualized daily standard deviation of returns over those 47 days. The underlying hypothesis is that if the forecast is unbiased the parameters α and β should be equal to 0 and 1 respectively. In evaluating which volatility measure is a better forecast of realized volatility I will then follow three criteria: which one has a smaller forecast error, in this case measured by the constant; which measure has a slope which closer is to one; and finally which forecasting model has the higher average adj-R^2.

Table 10 reports the estimation results. In the first panel I report results of the time series analysis wherein I estimate the coefficients of Equation 2 for each stock using the entire available sample. The analysis delivers a pair of estimated parameters $(\hat{\alpha}, \hat{\beta})_{i=1:N}$ for each stock. I tabulate the cross-sectional mean of the coefficients as well as the standard deviation of the cross-sectional distribution (in parenthesis). In the second panel I tabulate the results of the analysis when Equation 2 is estimated cross-sectionally. Following Fama and MacBeth (1973) a cross-sectional regression is estimated at each date, and then coefficients are averaged through time. Average parameters and t-statistics corrected for serial dependence are shown in the table. Invariably both panels offer the same evidence: CS is a better predictor of future realized volatility. In both the time-series and the cross-sectional approach model (1) has a smaller $\hat{\alpha}$, a larger $\hat{\beta}$ and a higher average adj-R^2. For example in the time series regression in 64% of the cases the constant from model (1) is closer to zero than the constant from model (2). The proportion of times the CS slope is closer to one is instead 65%. Finally in 73% of

[12]There are several articles in the finance literature study the relation between implied volatility and future realized volatility. Among many others we find Day and Lewis (1992), Canina and Figlewski (1993), Fleming (1998), Lamoureux and Lastrapes (1993), and Jiang and Tian (2005).

the cases the adjusted-R^2 is greater.[13] In the Fama-MacBeth regression the estimated intercept in model (1) is very small, –0.007, and not statistically different from zero, while the slope estimated coefficient is 1.011 and not statistically different from 1. In this context inference about the differences in the estimated coefficients form across models is directly possible. The mean difference of the constants (–0.035) is statistically significant from zero with a t-statistic, corrected for serial dependence, of –9.7. Similarly the mean difference in the slopes (0.080) has a t-statistic of 12.6.

Having shown that the cross-sectional forecast CS is a better prediction of future realized volatility, I investigate if it leads to informationally efficient option prices. In particular, I price the options involved in the portfolio strategies previously discussed by plugging the CS estimate into the Black and Scholes (1973) model and show that the trading strategy is no longer profitable. For each of the decile groups obtained by sorting the out-of-sample volatility forecasts, I compute the portfolio returns using as initial prices the values obtained by inserting CS into the Black and Scholes formula. I use the LIBOR rate as the interest rate, while the dividend yield is calculated from the last dividend paid by the firm. Summary statistics, average and t-statistic, are reported in Table 11. Note that the ten portfolio average returns are not statistically different form zero; the average return of portfolio (1) goes from –6.7% to 2.2% while the return of portfolio (10) goes from 8.4% to 1.8%. The result suggests that the cross-sectional forecast CS works better for the high portfolios, for which even the magnitude of the average returns is very close to zero. This suggests that the modified price obtained from the volatility forecast is closer to the 'true' price of the option.

6 Conclusion

In this paper I document the existence of predictability in the cross-sectional distribution of equity option implied volatilities. Various implications of the efficient market hypothesis have already been examined in context of option markets. However, differently from these papers, I document that the predicatibility found in this paper is economically significant to investors.

[13]In general, standard inference on the cross-sectional distribution of estimated parameters is not easily available, because the estimates are likely to be cross-correlated. If we assume that the estimates are not cross-correlated standard inference applies: both the difference in the constant from model (1) and (2), 0.023, and the difference in the slope, 0.052, are highly statistically significant with t-statistics of –7.1 and 7.9, respectively.

One possibility is that the cross-sectional predictability stems from the fact that economic agents do not use all available information in forming expectations about future stock volatilities. I show that a better prediction of the future realized volatility can be constructed using the cross-sectional forecasting model.

The verdict about what generates this behavior on the part of the economic agents is left for future research. In particular two possibilities appear as leading candidates: behavioral biases in the process in which information is incorporated into option prices, as suggested by Poteshman (2001) and Stein (1989) for example, or the existence of premia in the equity option market for exposure to aggregate risks such as volatility, Ang, Hodrick, Xing, and Xiaoyan (2005), and liquidity, Pedersen and Acharya (2005).

References

Amin, Kaushik, Joshua D. Coval, and Nejat H. Seyhun, 2004, Index option prices and stock market momentum, *Journal of Business* 77, 835–873.

Andersen, Torben G., Tim Bollerslev, Peter F. Christoffersen, and Francis X. Diebold, 2005, Volatility and Correlation Forecasting, forthcoming in *Handbook of Economic Forecasting*.

Ang, Andrew, Robert Hodrick, Yuhang Xing, and Zhing Xiaoyan, 2005, The cross-section of volatility and expected returns, forthcoming in *Journal of Finance*.

Bakshi, Gurdip, Nikunj Kapadia, and Dilip Madam, 2003, Stock returns characteristics, skew laws, and the differential pricing of individual equity options, *Review of Financial Studies* 16, 101–143.

Black, Fischer, 1976, Studies of stock price volatility changes, *Journal of tha American Statistical Association*, 177–181.

Black, Fischer, and Myron Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–654.

Canina, Linda, and Stephen Figlewski, 1993, The informational content of implied volatility, *Review of Financial Studies* 6, 659–681.

Carhart, Mark M., 1997, On persistence in mutual fund perfromance, *Journal of Finance* 52, 57–82.

Chan, Kalok, Peter Y. Chung, and Herb Johnson, 1995, The intraday behavior of bid-ask spreads for NYSE stocks and CBOE options, *Journal of Financial and Quantitative Analysis* 30, 329–346.

Christensen, Bent J., and Nagpurnanand R. Prabhala, 1998, The relation between implied and realized volatility, *Journal of Financial Economics* 50, 125–150.

Cox, John, Stephen Ross, and Mark Rubinstein, 1979, Option pricing: a simplified approach, *Journal of Financial Economics* 7, 229–263.

Das, Sanjiv Ranjan, and Rangarajan K. Sundaram, 1999, Of smiles and smirks: a term structure perspective, *Journal of Financial and Quantitative Analysis* 34, 211–239.

Day, Theodore E., and Craig M. Lewis, 1992, Stock market volatility and the information content of stock index options, *Journal of Econometrics* 52, 267–287.

De Fontnouvelle, Patrick, Raymond P.H. Fisher, and Jeffrey H. Harris, 2003, The behavior of bid-ask spreads and volume in options markets during the competition for listings in 1999, *Journal of Finance* 58, 2437–2463.

Dennis, Patrick, and Stewart Mayhew, 2002, Risk-neutral skewness: evidence from stock options, *Journal of Financial and Quantitative Analysis* 37, 471–493.

Dennis, Patrick, Stewart Mayhew, and Chris Stivers, 2005, Stock returns, implied volatility innovations, and the asymmetric volatility phenomenon, forthcoming in *Journal of Financial and Quantitative Analysis*.

Dubisnky, Andrew, and Michael Johannes, 2005, Earnings announcements and equity options, Working paper.

Easley, David, Maureen O'Hara, and P.S. Srinivas, 1998, Option volume and stock prices: evidence on where informed traders trade, *Journal of Finance* 53, 431–465.

Efron, Bradley, and Carl Morris, 1977, Stein's paradox in statistics, *Scientific American* 236, 119–127.

Fama, Eugene F., 1998, Market efficiency, long-term returns, and behavioral finance, *Journal of Financial Economics* 49, 283–306.

Fama, Eugene F., and Kenneth R. French, 1992, The cross-section of expected stock returns, *Journal of Finance* 47, 427–465.

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.

Fama, Eugene F., and James D. MacBeth, 1973, Risk, return, and equilibrium: empirical tests, *Journal of Political Economy* 81, 607–636.

Figlewski, Stephen, 1989, Options arbitrage in imperfect markets, *Journal of Finance* 44, 1289–1311.

Fleming, Jeff, 1998, The quality of market volatility forecasts implied by S&P 100 index option prices, *Journal of Empirical Finance* 5, 317–345.

George, Thomas J., and Francis A. Longstaff, 1993, Bid-ask spreads and trading activity in the S&P 100 index option market, *Journal of Financial and Quantitative Analysis* 28, 381–397.

Gould, John P., and Dan Galai, 1974, Transactions costs and the relationship between put and call prices, *Journal of Financial Economics* 1, 105–129.

Granger, Clive W.J., and Ser-Huang Poon, 2003, Forecasting volatility in financial markets: a review, *Journal of Economic Literature* 41, 478–539.

Ho, Thomas S.Y. snd Macris, Richard G., 1984, Dealer bid-ask quotes and transaction prices: an empirical study of some AMEX options, *Journal of Finance* 39, 23–45.

Jegadeesh, Narasimhan, 1990, Evidence of predictable behavior of security returns, *Journal of Finance* 45, 881–898.

Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to buying winners and selling losers: implications for stock market efficiency, *Journal of Finance* 47, 65–91.

Jiang, George J., and Yisong S. Tian, 2005, The model-free implied volatility and its information content, *Review of Financial Studies* 18, 1305–1342.

Karolyi, Andrew G., 1992, Predicting risk: some new generalizations, *Management Science* 38, 57–74.

Lamoureux, Christopher G., and William D. Lastrapes, 1993, Forecasting stock-return variance: toward an understanding of stochastic implied volatilities, *Review of Financial Studies* 6, 293–326.

Ledoit, Olivier, and Michael Wolf, 2003, Improved estimation of the covariance matrix of stock returns with an application to portfolio selection, *Journal of Empirical Finance* 10, 603–621.

Mayhew, Stewart, 2002, Competition, market structure, and bid-ask spreads in stock option markets, *Journal of Finance* pp. 931–958.

Morris, Carl, 1983, Parametric empirical bayes inference: theory and application, *Journal of tha American Statistical Association* 236, 119–127.

Ofek, Eli, Matthew Richardson, and Robert F. Whitelaw, 2004, Limited arbitrage and short sales restrictions: evidence from the options market, *Journal of Financial Economics* 74, 305–342.

Pan, Jun, and Allen M. Poteshman, 2005, The information in option volume for future stock prices, forthcoming in *Review of Financial Studies.*

Pedersen, Lasse H., and Viral Acharya, 2005, Asset pricing with liquidity risk, *Journal of Financial Economics* 77, 375–410.

Poteshman, Allen M., 2001, Underreaction, overreaction, and increasing misreaction to information in the options market, *Journal of Finance* 56, 851–876.

Roll, Richard, 1984, A simple implicit measure of the effective bid-ask spread in an efficient market, *Journal of Finance* 4, 1127–1139.

Santa-Clara, Pedro, and Alessio Saretto, 2005, Option strategies: good delas and margin calls, UCLA working paper.

Stein, Charles, 1955, Inadmissibility of the usual estimator for the mean of a multivariate normal distribution, *Proceedings of the Third Berkeley Symposium* 1, 197–206.

Stein, Jeremy, 1989, Overreactions in the options market, *Journal of Finance* 44, 1011–1023.

Swidler, Steve, and David J. Diltz, 1992, Implied volatilities and transaction costs, *Journal of Financial and Quantitative Analysis* 27, 437–447.

Toft, Klaus Bjerre, and Brian Prucyk, 1997, Options on leverraged equity: theory and empirical tests, *Journal of Finance* 52, 1151–1180.

Figure 1: VIX vs. IMP

In this figure I plot the time-series of VIX and the time-series of the average implied volatility, IMP. VIX data is obtained from the CBOE. Options and stocks closing prices were sampled monthly between January 1996 and June 2005. The data is provided by the Ivy DB database from OptionMetrics. All options are American.

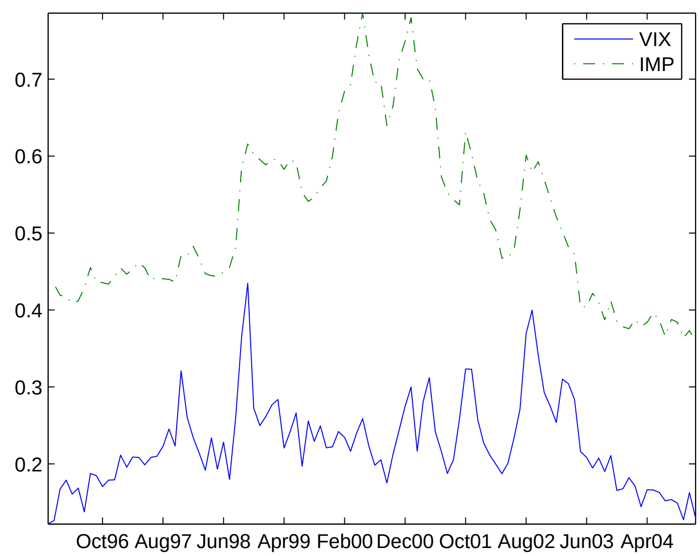

Figure 2: Emprical Distribution of CS

In Panel A I plot ΔIV^* versus the implied volatility as of the beginning of September 2001. In Panel B I plot the tails of the non-parametric kernel density of the cross-sectional distribution of IV and CS, where

$$CS_t = IV_t + \underbrace{(IV_t \times \Delta iv_t^*)}_{\Delta IV_t^*}$$

The data is provided by the Ivy DB database from OptionMetrics. All options are American.

Panel A

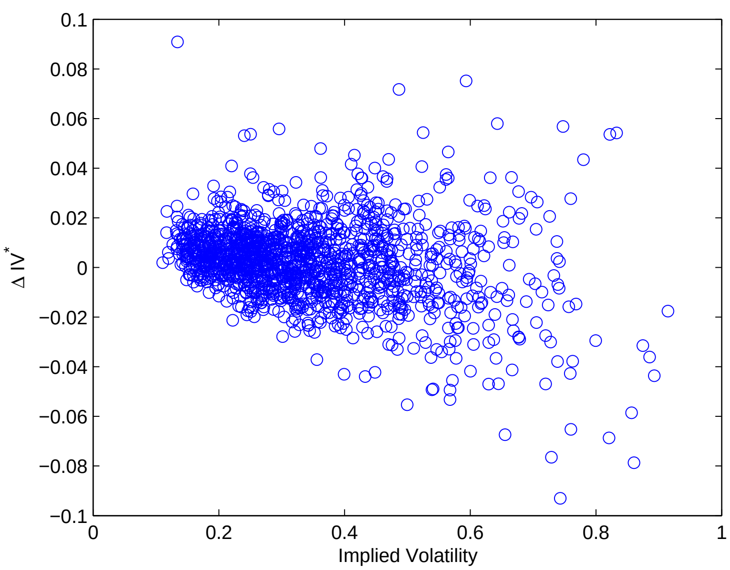

Panel B

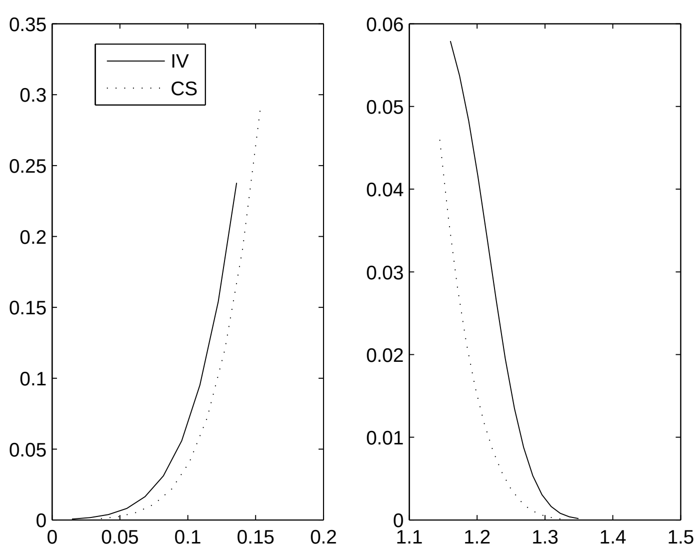

Table 1: Summary Statistics

This table reports summary statistics of call, put, and straddle returns as well as the level and change of the ATM implied volatilities, the smiles, the level and change of the realized volatilities. I report average, standard deviation, minimum, maximum, skewness, kurtosis. Options and stocks closing prices were sampled monthly between January 1996 and June 2005. The data is provided by the Ivy DB database from OptionMetrics. All options are American.

	mean	median	std	min	max	skew	kurt
Call	0.034	-0.370	1.185	-0.995	26.097	2.976	24.649
Put	-0.036	-0.400	1.065	-0.994	23.857	2.565	16.334
Strad	-0.003	-0.176	0.524	-0.896	12.030	2.912	23.045
IV	0.503	0.447	0.239	0.124	2.019	1.162	4.484
ΔIV	-0.002	-0.003	0.087	-0.868	0.774	0.334	10.589
SmL	0.031	0.002	0.054	-0.494	0.706	1.675	8.600
SmR	-0.005	0.000	0.035	-0.413	0.614	1.312	17.749
RV	0.491	0.410	0.307	0.029	3.510	1.969	9.313

Table 2: Volatility Predictability

In this table I report results of the estimation of the following forecasting model for the change in implied volatility:

$$\begin{aligned}\Delta iv_{i,t+1} &= \alpha_t + \beta_t iv_{i,t} + \gamma_t X_{i,t} + \epsilon_{t+1} \\ X_{i,t} &= \{iv_{i,t} - \overline{iv}_{i,t-12:t-1}, iv_{i,t} - \overline{rv}_{i,t-12:t-1}\}\end{aligned}$$

where $IV_{i,t}$ is the ATM implied volatility for stock i measured at the beginning of month t, $\Delta iv_{i,t} = \log(IV_{i,t+1}) - \log(IV_{i,t})$, $\overline{iv}_{i,t-12:t-1}$ is the natural logarithm of the twelve months moving average of IV_i and $\overline{rv}_{i,t-12:t-1}$ is the logarithm of the twelve months moving average of the realized volatility for stock i, $RV_{i,t-1}$, measured as the standard deviation of daily returns realized during month $t-1$. Averages of the cross-sectional estimates as well as Fama-MacBeth t-statistics adjusted for serial dependency are reported. Options and stocks closing prices were sampled monthly between January 1996 and June 2005. The data is provided by the Ivy DB database from OptionMetrics. All options are American.

	(1)	(2)	(3)	(4)
iv_t	-0.069	-0.038	-0.057	-0.044
	[-12.27]	[-5.90]	[-10.07]	[-7.18]
$iv_t - \overline{iv}_{t-12,t-1}$		-0.290		-0.191
		[-34.08]		[-16.66]
$iv_t - \overline{rv}_{t-12,t-1}$			-0.234	-0.121
			[-22.72]	[-10.79]
R^2	0.048	0.154	0.144	0.171
	{0.05}	{0.07}	{0.07}	{0.07}
$RMSE \times 100$				
mean	15.287	14.378	14.478	14.271
median	14.279	13.381	13.485	13.325
std	{8.57}	{8.04}	{7.98}	{7.96}
obs	398	398	398	398
	{94}	{94}	{94}	{91}

Table 3: Forecast and Future Implied Volatility

In this table, for each of the decile groups obtained by sorting the out-of-sample volatility forecasts (predicted logarithmic changes), I report the mean and the t-statistic for the difference between the level of IV and the cross-sectional mean of implied volatility, $\hat{IV}$, and the difference between IV and the twelve months moving average of past implied and realized volatilities, $\overline{IV}$. In the second panel I report the mean and the t-statistic of the out-of-sample predicted change (in levels), ΔIV^* as well as of the actual change, ΔIV. Options and stocks closing prices were sampled monthly between January 1996 and June 2005. The data is provided by the Ivy DB database from OptionMetrics. All options are American.

Δiv^*	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$IV_t - \hat{IV}_t$	0.112	0.065	0.034	0.010	-0.005	-0.025	-0.037	-0.053	-0.064	-0.080
t-stat	[13.27]	[9.68]	[6.24]	[2.28]	[-1.66]	[-7.71]	[-8.43]	[-8.28]	[-7.67]	[-8.00]
$IV_t - \overline{IV}_{t-12:t-1}$	0.150	0.060	0.027	0.006	-0.011	-0.022	-0.037	-0.048	-0.066	-0.118
t-stat	[15.42]	[8.04]	[4.07]	[0.93]	[-1.91]	[-3.91]	[-6.88]	[-8.96]	[-12.51]	[-21.63]
ΔIV^*_{t+1}	-0.073	-0.034	-0.019	-0.010	-0.002	0.004	0.010	0.016	0.024	0.039
t-stat	[-19.04]	[-11.31]	[-7.05]	[-3.72]	[-0.89]	[1.83]	[4.23]	[6.61]	[9.05]	[13.47]
ΔIV_{t+1}	-0.059	-0.026	-0.013	-0.003	0.000	0.006	0.010	0.014	0.020	0.035
t-stat	[-11.39]	[-5.52]	[-3.13]	[-0.85]	[0.07]	[1.56]	[2.70]	[4.45]	[5.99]	[10.29]

Table 4: Conditional Portfolio Returns

In this table,for each of the decile groups obtained by sorting the out-of-sample volatility forecasts (predicted percentage changes), I report summary statistic of the portfolios of call, put, and straddle returns: mean, standard deviation, minimum, maximum, Sharpe ratio (SR) and certainty equivalent (CE). CE is computed from a utility function with constant relative risk aversion parameter of 3. In the last column I report results for a zero-cost portfolio which is long in the options with the highest predicted increase in volatility and short in the highest decreases. Options and stocks closing prices were sampled monthly between January 1996 and June 2005. The data is provided by the Ivy DB database from OptionMetrics. All options are American.

Δiv^*	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(10-1)
CALL											
mean	-0.031	0.033	0.017	0.023	0.040	0.054	0.042	0.039	0.051	0.088	0.118
std	0.415	0.474	0.438	0.448	0.448	0.453	0.478	0.452	0.451	0.486	0.354
min	-0.801	-0.811	-0.744	-0.849	-0.795	-0.708	-0.750	-0.825	-0.791	-0.752	-1.071
max	1.237	1.125	1.441	1.228	1.106	1.288	1.400	1.464	1.165	1.617	1.252
SR	-0.081	0.062	0.032	0.045	0.083	0.112	0.081	0.079	0.106	0.174	0.334
CE	-0.301	-0.306	-0.278	-0.316	-0.286	-0.239	-0.284	-0.275	-0.271	-0.223	-0.651
PUT											
mean	-0.120	-0.084	-0.060	-0.052	-0.038	-0.049	-0.012	-0.006	0.002	0.056	0.176
std	0.418	0.474	0.479	0.494	0.496	0.493	0.499	0.514	0.514	0.537	0.291
min	-0.760	-0.737	-0.672	-0.716	-0.730	-0.753	-0.700	-0.762	-0.706	-0.684	-0.503
max	1.825	2.182	2.586	2.709	2.626	2.486	2.454	2.918	3.059	2.959	1.134
SR	-0.294	-0.183	-0.132	-0.110	-0.084	-0.105	-0.030	-0.018	-0.002	0.099	0.605
CE	-0.358	-0.376	-0.313	-0.320	-0.317	-0.322	-0.286	-0.299	-0.265	-0.232	0.068
STRADDLE											
mean	-0.072	-0.023	-0.024	-0.013	-0.003	0.004	0.011	0.009	0.019	0.056	0.129
std	0.123	0.139	0.148	0.146	0.145	0.147	0.156	0.159	0.157	0.175	0.132
min	-0.315	-0.267	-0.280	-0.268	-0.254	-0.290	-0.234	-0.245	-0.288	-0.196	-0.216
max	0.491	0.561	0.722	0.754	0.778	0.700	0.808	0.895	0.969	0.927	0.486
SR	-0.616	-0.187	-0.184	-0.107	-0.038	0.006	0.050	0.037	0.103	0.304	0.977
CE	-0.093	-0.049	-0.051	-0.039	-0.028	-0.022	-0.018	-0.020	-0.010	0.022	0.105
STOCK											
mean	0.008	0.012	0.011	0.012	0.010	0.013	0.008	0.007	0.009	0.007	-0.001
std	0.086	0.088	0.078	0.074	0.071	0.067	0.066	0.058	0.059	0.056	0.062
min	-0.245	-0.237	-0.209	-0.227	-0.235	-0.198	-0.196	-0.194	-0.201	-0.181	-0.265
max	0.227	0.193	0.276	0.263	0.187	0.170	0.120	0.115	0.121	0.140	0.178
SR	0.056	0.101	0.099	0.121	0.099	0.145	0.079	0.077	0.102	0.076	-0.009
CE	-0.004	-0.000	0.002	0.004	0.002	0.006	0.001	0.002	0.004	0.002	-0.007

Table 5: Straddle Portfolio Returns in Subsamples

In this table I report summary statistic for decile portfolios of straddle returns: mean and t-statistic. In the last column I report results for a zero-cost portfolio which is long in the options with the highest predicted increase in volatility and short in the highest decreases. I report results for different states of the stock market, up and down market, and for periods of increasing and decreasing aggregate volatility. A test on the difference of the means in also reported. The market is proxied by the value-weighted CRSP market portfolio while aggregate volatility is measured by VIX. Options and stocks closing prices were sampled monthly between January 1996 and June 2005. The data is provided by the Ivy DB database from OptionMetrics. All options are American.

Δiv^*	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(10-1)
						up market					
mean	-0.059	-0.010	-0.018	-0.011	0.001	0.004	0.011	0.001	0.012	0.048	0.108
t-stat	[-4.15]	[-0.59]	[-1.01]	[-0.66]	[0.04]	[0.23]	[0.63]	[0.06]	[0.74]	[2.62]	[6.32]
						down market					
mean	-0.087	-0.039	-0.032	-0.012	-0.005	0.006	0.013	0.026	0.034	0.070	0.156
t-stat	[-4.18]	[-1.79]	[-1.27]	[-0.49]	[-0.18]	[0.26]	[0.47]	[0.92]	[1.17]	[2.13]	[8.06]
						up - down					
mean	0.027	0.029	0.014	0.001	0.005	-0.002	-0.001	-0.025	-0.022	-0.021	-0.048
t-stat	[1.08]	[1.03]	[0.47]	[0.04]	[0.17]	[-0.07]	[-0.03]	[-0.75]	[-0.65]	[-0.56]	[-1.87]
						up volatility					
mean	-0.073	-0.014	-0.016	0.012	0.025	0.038	0.045	0.052	0.062	0.104	0.176
t-stat	[-3.61]	[-0.64]	[-0.71]	[0.54]	[1.09]	[1.67]	[1.85]	[2.02]	[2.50]	[3.44]	[9.00]
						down volatility					
mean	-0.073	-0.033	-0.033	-0.037	-0.026	-0.027	-0.021	-0.029	-0.018	0.012	0.085
t-stat	[-5.39]	[-1.87]	[-1.79]	[-2.09]	[-1.58]	[-1.53]	[-1.15]	[-1.73]	[-1.05]	[0.76]	[5.67]
						up - down					
mean	-0.000	0.019	0.017	0.049	0.052	0.065	0.066	0.081	0.080	0.092	0.092
t-stat	[-0.00]	[0.70]	[0.57]	[1.70]	[1.81]	[2.26]	[2.17]	[2.63]	[2.64]	[2.69]	[3.72]

Table 6: Straddle Returns by Volatility Forecast and Stock Characteristics

In this table I report the mean returns and t-statistics of quintile straddle portfolios obtained by independently sorting stocks based on the forecast of the log change in volatility and risk-based firm characteristics: beta, market size, and book-to-market ratio. Options and stocks closing prices were sampled monthly between January 1996 and June 2005. The data is provided by the Ivy DB database from OptionMetrics. All options are American.

Δiv^*	L=1	2	3	4	H=5	H-L	L=1	2	3	4	H=5	H-L
BETA			*mean*						*t-stat*			
L=1	-0.054	-0.024	-0.003	0.001	0.047	0.101	-3.42	-1.14	-0.17	0.07	2.67	6.05
2	-0.060	-0.040	-0.013	-0.023	0.018	0.078	-4.33	-2.54	-0.80	-1.33	1.01	4.78
3	-0.051	-0.051	-0.032	0.005	0.023	0.074	-3.18	-3.24	-2.05	0.28	1.19	4.51
4	-0.055	-0.018	0.007	0.010	0.033	0.088	-3.38	-1.04	0.41	0.52	1.56	4.76
H=5	-0.031	0.021	0.036	0.058	0.085	0.116	-1.76	0.87	1.59	2.58	3.32	5.17
H-L	0.023	0.044	0.039	0.057	0.038		1.27	1.85	1.64	2.55	1.41	
SIZE			*mean*						*t-stat*			
S=1	-0.038	-0.006	0.003	0.043	0.088	0.126	-2.03	-0.21	0.12	1.41	2.37	3.23
2	-0.030	0.007	0.023	0.058	0.066	0.096	-2.18	0.35	1.11	2.42	3.17	4.79
3	-0.035	-0.011	0.035	0.035	0.059	0.094	-2.42	-0.72	2.24	1.82	3.29	7.20
4	-0.049	-0.003	0.006	0.022	0.039	0.088	-3.30	-0.20	0.36	1.21	2.32	5.85
B=5	-0.077	-0.050	-0.026	-0.025	0.012	0.089	-5.34	-3.54	-1.69	-1.60	0.70	6.53
S-B	0.039	0.044	0.029	0.068	0.076		1.87	1.70	1.17	2.19	2.23	
BM			*mean*						*t-stat*			
L=1	-0.053	-0.015	-0.002	0.004	0.034	0.087	-3.94	-1.03	-0.11	0.26	2.33	6.69
2	-0.051	-0.043	-0.010	0.013	0.036	0.087	-3.26	-2.96	-0.62	0.73	2.14	6.63
3	-0.056	-0.018	0.016	0.024	0.043	0.100	-3.81	-0.91	0.98	1.17	2.15	5.77
4	-0.062	-0.024	-0.026	0.018	0.066	0.128	-3.95	-1.27	-1.32	0.84	2.70	5.91
H=5	-0.046	0.016	0.041	0.052	0.054	0.101	-2.28	0.65	1.50	1.66	2.02	3.55
H-L	0.007	0.032	0.042	0.048	0.020		0.33	1.33	1.57	1.68	0.79	
MOM			*mean*						*t-stat*			
D=1	-0.034	-0.002	0.035	0.036	0.074	0.107	-2.19	-0.15	1.72	1.70	3.15	4.82
2	-0.057	-0.010	-0.006	0.010	0.048	0.105	-4.13	-0.65	-0.34	0.58	2.64	6.75
3	-0.053	-0.044	-0.018	0.005	0.026	0.079	-3.35	-2.93	-1.27	0.29	1.47	4.11
4	-0.047	-0.025	0.017	-0.007	0.022	0.069	-2.95	-1.55	0.79	-0.41	1.26	4.30
U=5	-0.055	-0.028	0.015	0.022	0.033	0.088	-3.07	-1.46	0.84	1.04	1.86	4.80
U-D	-0.021	-0.026	-0.020	-0.014	-0.041		-1.12	-1.43	-0.86	-0.64	-1.78	

Table 7: Risk-Adjusted Straddle Returns

In this table I report the estimation results from regressing the straddle portfolio returns on a linear pricing model composed by the Fama and French three factors model, the Carhart momentum factor, and changes in the aggregate implied volatility, as measured by VIX. I run the regression for portfolio (1) (lowest predicted logarithmic increase in volatility), (10) (highest predicted logarithmic increase in volatility) and (10-1). Options and stocks closing prices were sampled monthly between January 1996 and June 2005. The data is provided by the Ivy DB database from OptionMetrics. All options are American.

Δiv^*		(1)			(10)			(10-1)	
$const$	-0.068	-0.075	-0.069	0.073	0.053	0.071	0.141	0.129	0.140
	[-3.94]	[-5.78]	[-4.07]	[3.24]	[2.95]	[3.42]	[11.98]	[9.68]	[12.82]
MKT	-0.465		-0.355	-1.209		-0.737	-0.745		-0.382
	[-0.82]		[-0.70]	[-1.74]		[-1.18]	[-2.75]		[-1.33]
SMB	0.195		0.216	-1.341		-1.250	-1.536		-1.466
	[0.50]		[0.56]	[-2.70]		[-2.56]	[-4.74]		[-4.63]
HML	-0.700		-0.704	-1.971		-1.987	-1.271		-1.283
	[-1.32]		[-1.35]	[-3.26]		[-3.54]	[-3.66]		[-3.63]
MOM	-0.195		-0.219	0.079		-0.025	0.274		0.194
	[-1.05]		[-1.14]	[0.31]		[-0.09]	[1.35]		[0.88]
ΔVIX		0.166	0.280		1.259	1.209		1.093	0.929
		[0.38]	[0.85]		[2.02]	[2.81]		[3.59]	[3.65]
$adj - R^2$	0.019	-0.007	0.015	0.130	0.076	0.179	0.222	0.104	0.274

Table 8: Straddle Holding Period Returns at Different Horizons

In this table I report summary statistic for decile portfolios of straddle returns constructed using forecasts over 1,3 and 6 month horizon. I tabulate mean, standard deviation, minimum, maximum, Sharpe ratio (SR) and certainty equivalent (CE). CE is computed from a utility function with constant relative risk aversion parameter of 3. In the last column I report results for a zero-cost portfolio which is long in the options with the highest predicted increase in volatility and short in the highest decreases. Options and stocks closing prices were sampled monthly between January 1996 and June 2005. The data is provided by the Ivy DB database from OptionMetrics. All options are American.

	1 month	*3 month*	*6 month*
		maturity	
mean	0.129	0.059	0.036
std	0.132	0.069	0.073
min	-0.216	-0.138	-0.178
max	0.486	0.222	0.237
SR	0.977	0.843	0.494
CE	0.105	0.051	0.028
		horizon	
mean	0.129	0.149	0.150
std	0.132	0.196	0.319
min	-0.216	-0.575	-1.115
max	0.486	0.623	0.944
SR	0.977	0.762	0.471
CE	0.105	0.078	-0.347

Table 9: Impact of Liquidity and Transaction Costs

In this Table I report average returns and t-statistics for the long-short portfolio when the sample is split in different liquidity groups which are obtained by ranking stocks on the base of the liquidity characteristics of the options. For each stock I compute the average quoted bid-ask spread of all the options series traded in the previous month as well as the daily average dollar volume. I the sort stocks based on these options liquidity characteristics. I report the average return computed from the mid-point price (MidP) and from the effective bid-ask spread (EFP), estimated to be equal to 75% of the quoted spread. The same set of results is also tabulated for the case when the portfolio formation is lagged by one day (skip 1 day). Options and stocks closing prices were sampled monthly between January 1996 and June 2005. The data is provided by the Ivy DB database from OptionMetrics. All options are American.

(10–1)	no skip		skip 1 day	
	MidP	EBA	MidP	EBA
BA Spread				
low	0.136	0.072	0.128	0.065
	[3.68]	[1.95]	[3.54]	[1.81]
med	0.183	0.086	0.170	0.074
	[4.98]	[2.43]	[4.71]	[2.12]
high	0.166	0.029	0.154	0.017
	[4.91]	[0.91]	[4.65]	[0.54]
	MidP	EBA	MidP	EBA
Volume				
low	0.141	0.014	0.133	0.005
	[4.32]	[0.45]	[4.11]	[0.16]
med	0.184	0.088	0.169	0.073
	[5.55]	[2.77]	[5.17]	[2.34]
high	0.163	0.086	0.154	0.078
	[4.03]	[2.16]	[3.85]	[2.00]

Table 10: Forecasts of Future Realized Volatility

In this Table I report the estimation results of different forecasting models for the future realized volatility. The dependent variable is the realized volatility of daily returns over the life of the option. So, for example, if the implied volatility is extracted from a couple of options with 47 days to expiration I compute the annualized daily standard deviation of returns over those 47 days. The forecasting variables are the cross-sectional estimate (CS), the market implied volatility (IV), and the past realized volatility (RV). In the first panel I report results of a time series analysis wherein I run a forecasting regression for each stock; I tabulate the cross-sectional mean of the estimated coefficients as long as the standard deviation of the cross-sectional distribution. This value is tabulated in parenthesis. In the second panel I tabulate results of a Fama-MacBeth regression, t-statistics corrected for serial dependence are in brackets. Options and stocks closing prices were sampled monthly between January 1996 and June 2005. The data is provided by the Ivy DB database from OptionMetrics. All options are American.

	const	CS	IV	RV	*adj*-R^2
			time-series regression		
(1)	0.041	0.907			0.404
	(0.20)	(0.41)			
(2)	0.064		0.854		0.365
	(0.21)		(0.41)		
(3)	0.264			0.437	0.224
	(0.20)			(0.32)	
			cross-sectional regression		
(1)	-0.007	1.011			0.647
	[-0.68]	[42.43]			
(2)	0.028		0.931		0.628
	[2.48]		[36.38]		
(3)	0.125			0.725	0.522
	[11.50]			[33.45]	

Table 11: Pricing Options using CS

In this Table, for each of the decile groups obtained by sorting the out-of-sample volatility forecasts, I report summary statistic of the straddle portfolio returns. In the first Panel I tabulate the returns computed from the market prices (precisely as in Table 4). In the second Panel I report statistics for the portfolios when the returns are computed using as initial prices the values obtained by inserting the cross-sectional prediction of implied volatility (CS) into the Black and Scholes formula. I used the LIBOR rate as the interest rate, while the dividend yield is calculated based on the last dividend paid by the firm. t-statistics corrected for serial dependence are in brackets. Options and stocks closing prices were sampled monthly between January 1996 and June 2005. The data is provided by the Ivy DB database from OptionMetrics. All options are American.

Δiv^*	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
					market prices					
mean	-0.067	-0.038	-0.018	-0.017	-0.002	0.013	0.019	0.021	0.041	0.084
t-stat	[-5.46]	[-2.90]	[-1.27]	[-1.23]	[-0.16]	[0.83]	[1.18]	[1.31]	[2.41]	[4.23]
					prices based on CS					
mean	0.022	0.013	0.017	0.006	0.011	0.016	0.012	0.003	0.008	0.018
t-stat	[1.58]	[1.08]	[1.40]	[0.44]	[0.81]	[1.14]	[0.84]	[0.22]	[0.56]	[1.04]