Forthcoming in *Advances in Futures and Options Research*

THE ECONOMIC SIGNIFICANCE OF THE FORECAST BIAS OF S&P 100 INDEX OPTION IMPLIED VOLATILITY

Jeff Fleming[*]
Jones Graduation School of Management
Rice University

Abstract

A number of recent papers find that the volatility implied by index option prices significantly overstates future stock market volatility. We investigate whether this bias is purely due to measurement error and model misspecification, or whether the bias is also apparent in option market prices. We accomplish this by examining the profits for trading strategies designed to exploit the apparent bias. Ignoring transaction costs, the strategies consistently earn significant positive profits which indicates the bias is indeed a function of option prices. The degree of bias, however, does not signal market inefficiency because the profits disappear once we impose bid/ask transaction costs. Our analysis also reveals that the bias is too large to be explained by skewness preference, but that it may be the result of market imperfections (e.g., transaction costs) and/or a premium demanded for volatility risk. We also find that the bias apparent through the trading strategies emerged only after the 1987 stock market crash.

April 18, 1998

[*] Address correspondence to Jeff Fleming, Jones Graduate School – MS 531, Rice University, P.O. Box 1892, Houston, TX 77251-1892, phone 713/527-4677, e-mail jfleming@rice.edu. This paper has benefited from many useful comments and suggestions by Chris Kirby, Barbara Ostdiek, Peter Ritchken (the editor), and an anonymous referee.

I. Introduction

An option's value fundamentally depends on the volatility rate of the underlying asset. When the volatility rate is uncertain, the aggregation of volatility assessments across investors determines the market price. Therefore, if we can adequately specify the valuation model used by the market, we can invert the observed price and recover the market's expectation regarding future volatility. This implied volatility should represent the best possible volatility forecast because, if the option market is efficient, all relevant conditioning information has been collapsed into the option price.

A number of recent studies examine how well the implied volatility forecasts future volatility. For stock index options, the evidence is mixed. Canina and Figlewski (1993) find that the implied volatility contains little information regarding future volatility, while other studies, such as Feinstein (1989), Fleming, Ostdiek, and Whaley (1995), and Bates (1998), find a much stronger relation. Day and Lewis (1992) find that the implied volatility has explanatory power, but also that GARCH forecasts contain additional relevant information. Fleming (1998) finds that the implied volatility dominates historical volatility as a forecast, as well as many other parameters often used to generate volatility forecasts.

These studies suggest that the implied volatility of index options may be useful for forecasting stock market volatility. But there is also some consensus that it is a biased forecast. The studies by Fleming, Ostdiek, and Whaley, Bates, and Fleming, as well as Jorion (1995) using currency options, each specifically examines this issue and concludes that the implied volatility systematically overstates future volatility. Jackwerth and Rubinstein (1996) also provide indirect evidence to support this conclusion. This finding motivates an important question concerning the source of the implied volatility's bias. At one extreme, the bias suggests that the volatility rate used by the market is systematically incorrect, and therefore the option market may be inefficient. At the other extreme, however, the bias may be entirely an artifact of construction, caused by measurement error and/or misspecification of the valuation model used to recover the market's volatility forecast.

This study assesses the validity of these possibilities. Our first objective is to determine whether the bias is purely due to measurement error and model misspecification, or whether the bias is inherent in option market prices. We accomplish this by examining the profitability of trading strategies designed to exploit the apparent bias. The bias implies that option prices may be "too high," so to try and profit from this opportunity we sell options and delta-hedge the exposure. Such a strategy will earn abnormal profits only if the bias is truly a function of option market prices. If instead option prices are correct and the bias is caused by measurement error or misspecification of the model used in computing implied volatility, then the abnormal trading profits will be zero.

Our second objective, should the trading strategies yield positive profits, is to determine the source of the bias. Market inefficiency is one possibility, but even in an efficient market both option prices and the implied volatility can seem biased due to transaction costs and other market imperfections, investor preferences regarding skewness and higher-order moments of returns, or a premium for volatility and/or jump risk. By examining the nature of the trading strategy profits, we can assess the relative importance of these factors in contributing to the observed bias. This evidence, in turn, provides guidance for future research which seeks to improve upon the forecast quality of the implied volatility, and, more generally, to improve on the performance of option valuation models.

Our empirical results indicate that the trading strategies do earn large profits before we impose any transaction costs. The strategies earn Sharpe ratios that are three to five times as large as those for investing in the S&P 100 index, and the profits are systematic throughout the entire sample period. This finding indicates that the implied volatility's forecast bias is indeed a function of option prices. Once we account for bid/ask transaction costs, however, the profits disappear. Therefore, the bias does not signal inefficiency in the option market.

Our analysis of the trading strategy profits focuses on three issues. First, the distribution of profits is negatively skewed, so the mean/variance performance of the strategies may seem abnormal only because they sacrifice skewness. Using Leland's (1996) modified-beta approach, however, we find that the higher-order moments of returns explain only a small portion of the abnormal performance. Second, the strategies entail volatility risk, and their profitability may represent compensation for this exposure. Consistent with this view, we find that the profits are greater for the strategies with greater volatility risk and during periods when the level of volatility (i.e., the volatility risk premium) is higher. Finally, transaction costs and other market imperfections can allow option prices to deviate from their "true" values without signaling arbitrage opportunities. The structure of these deviations may have systematically changed after the 1987 stock market crash as this event dramatically increased institutional buying pressure in the index option market. In support of this argument, we find that the trading strategies became profitable only after the crash.

The remainder of this paper is organized as follows. Section II examines the theoretical relation between the implied volatility and the market's volatility forecast, which we use to identify the possible sources of bias. Section III describes the trading strategies designed to exploit the apparent bias. Section IV reports the trading strategy profits. Section V examines the potential explanations regarding the source of the profits. The final section provides a summary and conclusion.

II. The Implied Volatility as a Forecast

A. Theoretical Motivation

The implied volatility from the Black/Scholes (1973) model is frequently used in both research and practice as a measure of expected volatility. The Black/Scholes model, however, assumes a known and constant volatility rate. The very fact that we are interested in measuring expected volatility suggests volatility is uncertain and time-varies. In this section, we develop an interpretation of the implied volatility when volatility time-varies. Specifically, we show that if volatility and jump risk are not priced and the correlation between asset returns and volatility is zero, then the implied volatility for nearby, at-the-money options should approximately equal the expected average volatility over the life of the option.

Consider first the special case where volatility is a deterministic function of asset price and/or time. Following Merton (1973), the riskless hedge argument of the Black/Scholes model is preserved, and an option's value equals its Black/Scholes value but with the (constant) volatility argument replaced by the average volatility over the life of the option (time t to T),

$$s_{t;T} = \sqrt{\frac{1}{T-t}\int_t^T s_x^2 \, dx} \, . \qquad (1)$$

More generally, when volatility is stochastic or the return process includes jumps, the riskless hedge argument no longer holds and volatility (or jump) risk can influence option value. Under the assumption that this risk is nonsystematic or investors are otherwise indifferent toward this risk, then option value remains unchanged. The value also generally depends on the correlation between volatility and returns. Hull and White (1987) and Scott (1987) show that when we assume this correlation is zero the option's value equals

$$f_t = \mathrm{E}[\mathrm{BS}(s_{t;T})|\boldsymbol{F}_t], \tag{2}$$

the expected Black/Scholes value, BS(·), evaluated at the average volatility over the life of the option, $s_{t;T}$, and conditioned on all information available at time t, $\boldsymbol{F}_t$,. Bates (1996) develops a similar result for the Merton (1976) jump-diffusion model with mean-zero jumps.

Feinstein (1989) demonstrates that for near-expiration, at-the-money options, the Black/ Scholes formula is approximately linear in its volatility argument. Therefore,

$$\mathrm{E}[\mathrm{BS}(s_{t;T})|\boldsymbol{F}_t] \approx \mathrm{BS}(\mathrm{E}[s_{t;T}|\boldsymbol{F}_t]), \tag{3}$$

and

$$f_t \approx \mathrm{BS}(\mathrm{E}[s_{t;T}|\boldsymbol{F}_t]). \tag{4}$$

As a result, the implied volatility, $\bar{s}_{t;T} = \mathrm{BS}^{-1}(f_t)$, satisfies

$$\bar{s}_{t;T} \approx \mathrm{E}[s_{t;T}|\boldsymbol{F}_t], \tag{5}$$

which indicates it should be an unbiased forecast of the average volatility over the life of the option, and its forecast error should be orthogonal to the market's information set at time t.[1]

This interpretation of implied volatility is consistent with recent studies that examine its forecast performance. In particular, most studies estimate a regression of implied volatility on future realized volatility. Under equation (5), the intercept and slope coefficients, respectively, should equal zero and one, and the regression error should be orthogonal to any conditioning variables (such as historical volatility or ARCH forecasts) available at time t. Underlying this test is a joint hypothesis that the implied volatility equals the market's volatility forecast and that this forecast equals the true conditional expectation. As a result, evidence supporting rejection may indicate that the option market is inefficient (i.e., it uses the wrong volatility forecast) and/or that equation (5) is misspecified (i.e., we have incorrectly measured the market's forecast). The likelihood of the second possibility is influenced by both error in measuring the implied volatility and the validity of the assumptions underlying equation (5).

For stock index options, the assumptions seem questionable on several grounds. First, the assumption of zero correlation between asset returns and volatility contradicts the leverage

[1] Note that equation (3), though approximate, is also biased. Since the Black/Scholes function is strictly concave in volatility, Jensen's inequality implies that the Black/Scholes valuation (using the average volatility rate) always exceeds the stochastic volatility valuation. The magnitude of this effect on implied volatility, however, is generally trivial, and its direction is opposite that of the observed forecast bias.

effect described by Black (1976). French, Schwert, and Stambaugh (1987) document a strong negative relation between index returns and volatility, and this relation may induce the skewness in index option prices observed by Bates (1998). Second, it may seem unreasonable that volatility risk does not influence option value. The reason is that volatility and index returns are correlated, and the underlying index is a market portfolio that is correlated with aggregate consumption. Finally, equation (5) only explicitly applies to European-style options. Most of the empirical research concerning the implied volatility, is based on S&P 100 options which are American-style and include daily wildcard options. These factors, as well as measurement error in estimating the implied volatility, may invalidate equation (5) and explain why the S&P 100 implied volatility systematically overstates future volatility.

Our goal in this study is to determine whether this observed bias is entirely due to misspecification and measurement error, or whether the bias is also apparent in option market prices. The bias suggests that the volatility expectation used by the market is systematically too high, and this may indicate that the option market is inefficient. However, even in an efficient market, option prices can seem biased due to transaction costs and other market imperfections, investor preferences regarding skewness and higher-order moments of returns, or a premium demanded for volatility and/or jump risk. These factors can allow the market's risk-neutral volatility expectation to systematically overstate the true expected volatility, contributing to the observed bias in implied volatility. We want to determine whether such bias is apparent in option prices and, if so, which of these factors is responsible.

B. Implied Volatility Bias Versus Option Price Bias

One approach for identifying the source of the implied volatility bias would be to consider more complex option valuation models. Our analysis of equation (5) reveals a number of shortcomings of a simple Black/Scholes-type specification, and perhaps by explicitly modeling the volatility process, or the jump process, or the correlation between index returns and volatility, we can pinpoint the source of the bias. But there are two problems with this approach. First, it requires additional parameter estimation which reduces the precision of the volatility estimates. Second, *any* model we use remains open to the misspecification issue. If we find that volatility forecasts under the model are biased, we still cannot distinguish between market inefficiency and model misspecification. Bates (1998), for example, finds bias for S&P 500 futures options even after fitting a two-factor stochastic volatility model with jumps.

We use an alternative approach that is independent of a valuation model. The approach is based on an important distinction about the various sources of forecast bias. Suppose the only bias in equation (5) is due to market inefficiency, and that this causes the implied volatility to overstate future volatility. The implied volatility bias occurs in this case because the underlying option price is also overstated. Alternatively, suppose the market correctly values the option and the sole source of bias is misspecification, say, because the correlation between index returns and volatility is nonzero. Now, the bias occurs even though the option price is correct because the implied volatility misrepresents the market's "true" volatility forecast.

The key to distinguishing between these effects is to directly examine observed option prices. In the first case, the implied volatility bias occurs precisely because the option price embeds a misstatement of expected volatility by market participants. In the second case, the price does not embed a volatility misstatement, it only appears so because the implied volatility is a biased measure of the market's expectation. Using this distinction, bias induced by measurement error and model misspecification would appear only in the implied volatility, while the bias caused by market inefficiency, transaction costs, and risk premium effects would also be apparent in option prices. We include the risk premium effects in this second category because the market "adjusts" its true volatility expectation to account for the risk premium.[2]

III. Empirical Methodology

In this section, we develop our methodology for examining the source of the implied volatility's forecast bias. Our basic approach is to develop two trading strategies aimed at exploiting the apparent bias. This focuses the analysis directly on option prices and avoids any bias induced by transforming price into implied volatility. As a result, if the trading strategies generate positive profits, we can conclude that the bias is not purely the result of measurement error and model misspecification, but that the bias is a function of option market prices.

A. Data

The primary data used in this study are the S&P 100 option prices provided by the Chicago Board Options Exchange. These data consist of the entire history of time-stamped bid and ask quotes for each option series, and the contemporaneous S&P 100 index level is appended to each quote. Our analysis also uses S&P 500 futures prices. These data, provided by the

[2] In other words, the option price varies with the risk premium even if the expected volatility rate is unchanged.

Chicago Mercantile Exchange, include time-stamp prices for all transactions which differ in price from the preceding trade.

In order to compute a time-series of implied volatilities, we also require the riskless interest rate and the expected S&P 100 dividends over the life the option. To proxy for the riskless rate, we use the effective yield on the treasury bill whose maturity most closely matches the option expiration but is at least 30 days.[3] The effective yield is computed from the average of the 3:00 CST bid and ask discounts reported in the *Wall Street Journal*. To proxy for expected dividends, we use the actual dividends paid by S&P 100 firms.[4] This series was constructed by Harvey and Whaley (1992b) through June 1988. For the subsequent period, the dividends were obtained from the *Standard & Poor's 100 Information Bulletin*.

For most of our analysis, we use daily sampling (at 3:00 CST) over the period from January 1988 through December 1993. This sample is consistent with those used in prior research regarding the forecast performance of the implied volatility. We begin the sample after the 1987 stock market crash so we can separately evaluate its effect on our results.

B. Implied Volatility Forecast Bias

Before explaining our trading strategies, we should verify that the implied volatility is indeed a biased forecast using our data. We use the following procedure to construct the implied volatility series. Each day, we first identify the nearby expiration, at-the-money option series. To limit the effect of option expirations, we move to the second nearby contract once the nearby is within two weeks of expiration.[5] Next, as in Harvey and Whaley (1992a), we collect all of the transaction prices for this option series during a ten-minute window centered around the stock market close. Finally, we solve for the volatility rate (i.e., the implied volatility) that generates theoretical option values that most closely match the observed prices. We use the Fleming/Whaley (1994) valuation model which accounts for discrete index dividends and early exercise and wildcard opportunities,[6] and we use Whaley's (1986) nonlinear regression to estimate the implied volatility from the window of option prices.

---

[3] We require at least 30 days to maturity for the T-bill because shorter-term bills are substantially less liquid.

[4] This perfect foresight assumption is not unreasonable because the options used in this study average about 30 days to expiration. Most firms declare their dividends at least a month prior to payment.

[5] See Day and Lewis (1988).

[6] Note that option value increases with early exercise and wildcard opportunities. Therefore, if we ignore these features in computing the implied volatility from a given option price, the implied volatility's upward bias would be even greater than observed. With the Fleming/Whaley (1994) approach, it is straightforward to account for these features and no more parameter estimation is required than for the Black/Scholes model.

To measure the forecast bias, we compare each day's implied volatility with the volatility realized over the remaining life of the option. Our estimator for realized volatility is the square root of the sample variance of daily S&P 100 returns. Figure 1 summarizes the comparison using the implied volatility for calls. Panel A shows the implied volatility's forecast errors over time, and Panel B provides a histogram of the errors. On average, the implied volatility overstates realized volatility by about three percentage points (stated on an annualized basis). This evidence is consistent with both the direction and magnitude found in prior research. Moreover, though not shown in the figure, the bias for put options is slightly greater than for calls. This difference is consistent with Fleming (1998).

C. Trading Strategies

We now develop the trading strategies used to evaluate whether the bias is purely the result of model misspecification and measurement error, or whether it is also apparent in option market prices. The upward bias in implied volatility suggests that option prices may be overstated, so to trade on the bias we should sell options. In a general sense, profits occur when the asset price remains constant (i.e., less volatile) so it does not offset the premium collected from selling the options. We will separately consider the profits for call and put options.

To focus directly on the apparent volatility mispricing, we hedge the market risk of these short option positions. An option's delta is the derivative of its value with respect to the underlying index level. By holding delta shares of the index against the short option position, we create a riskless hedge. Under the Black/Scholes assumptions, we can continuously rebalance the stock position over time and as the index level changes, so the hedge portfolio is riskless and should earn a return equal to the riskless interest rate. If, however, the option is overpriced (volatility is too high), the return will be greater.

In practice, the Black/Scholes assumptions do not hold and the hedge portfolio will be imperfect. One alternative is to use a more complex valuation model to determine the hedge ratio. This may be inferior in practice, however, because it still entails specification error and requires additional parameter estimation. Dumas, Fleming, and Whaley (1998) examine this issue, considering a variety of functional forms for the volatility process in valuing S&P 500 options. Their results indicate that although these processes achieve a better in-sample fit of observed option prices, the hedge ratios based on these processes are less reliable out-of-sample than simply using the Black/Scholes model.

In light of these considerations, we use two different strategies to hedge the short option positions. The first is simply delta-neutral hedging using S&P 500 futures. We use S&P 500 futures because the costs of hedging via the component S&P 100 stocks are prohibitive.[7] With this replacement, the hedge incurs additional risk due to tracking error (price movements for the two indexes may differ) and basis changes (price movements for cash and futures may differ). The second strategy is to hedge the call option with the put (i.e., a short volatility spread). Because they are both at-the-money options, the call and put offset each other to a degree. The potential losses, however, are unbounded. Therefore, we also consider a modification of this strategy where we protect the position with deep out-of-the-money calls and puts (i.e., a hedged volatility spread, or a butterfly). Using any of these hedges, our trading profits will be risky. We account for this risk when we interpret the results.

To implement the trading strategies, we adopt the following procedure. On the first day of the sample, at 3:00 CST, we sell the nearby (but at least 15 days to expiration), at-the-money call and put options. These are the same options that generate the implied volatility bias in Figure 1. At the same time, depending on our hedging strategy, we either delta hedge the position using futures or we buy deep out-of-the-money (30 index points[8]) calls and puts. On each subsequent day, at 3:00, we re-evaluate the position. If the options are nearer than 15 days to expiration or a new option series is nearer the money, we liquidate the old position and establish a new one using the same procedure as on the first day. If we continue to hold the position, we simply make any necessary adjustments to the delta hedge. We then move to the next day and repeat this procedure through the remainder of the sample period.[9]

The procedure for rolling into a new option series rather than holding each position to expiration addresses three issues concerning the interpretation of the results. First, note that even if the options are mispriced, selling them will not *always* earn positive profits because the future volatility rate is uncertain. As a result, the analysis requires a larger number of independent observations than would be available if each position were held to expiration. Second,

---

[7] Hedging with futures instead of the index portfolio affects the required return on the riskless hedge. Now, the investment is negative because we collect the option premium and we pay nothing to enter the futures position. Therefore, the return should be negative and equal to the riskless interest rate.
[8] On some days, strike prices this far from the money are not available. When this occurs, we use the strike prices that are furthest from the money.
[9] The trading strategies ignore the possibility of early option exercise. This seems reasonable, however, because the options are always at-the-money and at least 15 days to expiration so the chance of early exercise is remote.

our intent is to focus on nearby, at-the-money options because they demonstrate the observed forecast bias in implied volatility. If we continued to hold an option after it no longer satisfied this selection rule, it would affect our use of the trading profits for inference regarding the implied volatility bias. Finally, the distribution of the trading strategy profits (measured daily) is relatively constant over time. This occurs because, whether we continue holding an existing position or we liquidate it to initiate a new position, the next day's profits are always determined by the nearby, at-the-money option.

Following this procedure, each day's return is an independent observation based on the same options that exhibit implied volatility bias. With so much trading, however, one might question whether we can still profit from any volatility mispricing. In other words, we sell overpriced options one day only to repurchase them later at a price that is also overstated. Although this is true, the profits should still be apparent because the overpricing is less upon liquidation. To see this, first note that Fleming (1998) demonstrates the implied volatility's forecast bias does not vary with time to expiration. Second, an option's vega, the derivative of value with respect to volatility, systematically decreases as time to expiration decreases. As a result, a given volatility misstatement has a larger price effect today (when we sell the options) than it will at some point in the future (when we liquidate them).

IV. Empirical Results

In this section, we examine the daily profits for our trading strategies. We begin by ignoring transaction costs. This analysis allows us to determine whether the forecast bias of the implied volatility is also apparent in option prices. If it is not, then we can conclude that the bias is purely the result of measurement error and/or misspecification. Next, we examine whether profits exist after imposing transaction costs. This allows us to determine whether the implied volatility bias signals abnormal profit opportunities.

A. Is the Implied Volatility Bias Apparent in Option Prices?

To measure the trading strategy profits in the absence of transaction costs, we assume that each day's trading occurs at the average of the bid and ask prices available at 3:00 PM. For the futures trades (for delta hedging), we use the last transaction price through 3:00. Our two trading strategies, selling options and selling volatility spreads, involve "negative" investment (i.e., we receive the option premia). To facilitate the comparison of profits across days and across strategies, the investment proceeds for each trade are standardized to equal $100. The daily profits can therefore be interpreted as percentage returns.

Figure 2 illustrates the cumulative profits for each trading strategy. Recall that the spirit of each strategy is to create a riskless position which captures volatility "mispricing." To the extent that the mispricing is zero and the position is perfectly hedged, the profits should be negative and equal to the riskless rate. In contrast to this hypothesis, the profits depicted in Figure 2 are positive and quite large. This finding is consistent with the evidence that the implied volatility overstates future volatility. In addition, the total profits for the put option strategy are slightly greater than those using calls. This ordering conforms with the relative magnitudes of the implied volatility bias for call and put options reported in Fleming (1998).

Recall that the profit curves in Figure 2 are driven by constant $100 "investments" (i.e., not compounding reinvested profits). As a result, the relatively constant slopes indicate that the profits are systematic. In general, only two exceptions occur. First, for the call and put option strategies, severe spikes occur in January 1988 and October 1989. These spikes reflect the sensitivity of the delta-hedged profits to the large market movements during these periods, and they illustrate the imprecision of delta hedging. Second, the profits for each strategy seem to tail off during a prolonged period in 1991. Other than these exceptions, the profits for each strategy are fairly stable throughout the sample. This indicates that the apparent bias is systematic and not isolated to a particular period.

To this point, we have not accounted for the risk associated with these strategies. Recall that controlling for risk was our motivation for hedging the various positions. As the variation in the profit curves in Figure 2 illustrates, these hedges are imperfect and each strategy exhibits substantial variation in profitability. We begin to account for this risk by computing the daily excess returns for each strategy. These returns are based on the effective rate earned on the T-bill nearest (but not less than) 30 days to maturity. For comparison, we also consider the daily excess returns earned on the S&P 100 index. Table 1 summarizes these returns.

The excess returns are qualitatively similar to the profits shown in Figure 2. They are highest for the put option strategy, followed by the volatility spread and the call option strategy. Examining the returns by year, we again see that the profits are relatively systematic. The volatility spread strategy, for example, earns an average excess return of more than 0.9% per day in every year except 1991. The call and put option strategies exhibit more variation.

In fact, there seems to be a slight relation between these returns and S&P 100 returns. The call (put) option returns have a weak tendency to be higher during years when index returns are higher (lower). Because this pattern is inversely related to the pattern for unhedged short calls and puts, it suggests that the positions are “overhedged” by delta hedging.

Across the entire sample, the average excess returns for the trading strategies are on the order of 1% per day (0.65% for the hedged volatility spread). These profits seem very large relative to the 0.03% mean daily return for the S&P 100, but the standard deviations of the means indicate that the option strategies entail substantially more risk. To assess the significance of the returns, we compute the *t*-statistics for each strategy. These statistics, also reported in Table 1, measure the validity of the null hypothesis that the mean excess return is zero. Using the entire sample, the lowest *t*-statistic for any of the option strategies is 3.99 (for call options). This evidence indicates strong rejection of the null.

Although the trading strategies earn significant positive profits, it is possible that they merely represent compensation for market risk. To consider this possibility, we examine the Sharpe (or reward/risk) ratios for the different strategies. Let $\bar{r}_i$ and s_i, respectively, denote the mean excess return and volatility rate for strategy *i*. The Sharpe (1966) ratio is defined as

$$\text{SR}_i = \bar{r}_i / s_i \, . \tag{6}$$

Underlying SR_i is a “zero-investment strategy” of investing in asset *i* and financing the investment with riskless bonds. Different levels of risk and return can be achieved by varying the investments in asset *i* and riskless bonds, but the reward/risk ratio for all such combinations is constant and equal to SR_i. Therefore, as Sharpe (1994) demonstrates, investments that earn relatively larger Sharpe ratios are said to be “superior.”[10]

We can use the *t*-statistics in Table 1 to develop a comparison of the Sharpe ratios. To see this, note that the *t*-statistic for a particular strategy, $t_i = \bar{r}_i / (s_i / \sqrt{N})$, where *N* is the sample size, is proportional to its Sharpe ratio, $\text{SR}_i = t_i / \sqrt{N}$. Therefore, the ratio of the *t*-statistics for any two strategies equals the ratio of their Sharpe ratios. Comparing the *t*-statistics in Table 1 reveals that the Sharpe ratios are much larger for the option strategies than for the S&P 100

[10] A potential criticism of this approach is that it considers only the mean and volatility, ignoring any differences in payoffs across states of nature that may affect asset value. For comparing the trading strategies to the S&P 100 index, however, these differences should be small because the two investments are driven by the same economic factors (e.g., return and volatility of the market portfolio, riskless interest rates). In the next section, we use a more sophisticated risk measurement approach which accounts for the skewness and higher-order moments of returns.

index. The short volatility spread, for example, earns a Sharpe ratio nearly five times larger than a long index position. The call option strategy fares the worst, but its Sharpe ratio is still about three times greater than for the index.

The Sharpe ratio comparisons are generally robust across the sample period, but with two exceptions. First, during 1991, only the call option strategy outperforms the S&P 100 index. Second, the put option strategy in 1989 fares worse than the S&P 100. This second result is primarily due to the increase in volatility during to the mini-crash in October 1989. Other than these exceptions, in every year of the sample, each of the trading strategies earns a greater Sharpe ratio than investing in the S&P 100 index.

Based on these results, it appears that the option trading strategies do earn abnormal profits in the absence of transaction costs. Since these strategies are designed to exploit the implied volatility's forecast bias, the trading profits indicate that the bias is a function of option market prices and not entirely the result of misspecification or systematic measurement error. In the next section, we examine whether the trading strategy profits are still apparent after imposing transaction costs.

B. Does the Implied Volatility Bias Signal Abnormal Profit Opportunities?

Our previous results indicate that the implied volatility bias is apparent in option prices, so one source of the bias may be market inefficiency. We assess this issue by examining whether we can earn abnormal profits (after transaction costs) by trading on the bias. We use the same trading strategies as before, but now we include the transaction costs for executing each trade. If the strategies still yield abnormal returns, we would conclude that the bias is due (at least in part) to market inefficiency.

Due to the difficulty of quantifying the total transaction costs, we focus on just one component of an investor's costs, the bid/ask spread. Some of the remaining components, such as exchange fees, would be easy to quantify, while commissions and market-impact costs depend on the trade size and/or are difficult to measure. Instead of trying to pinpoint each of these costs, we simply rely on the bid/ask spread for a rough indication and note that the costs would be even higher if we include the other components.

Imposing bid/ask costs on the option trades is relatively straightforward. We simply assume that each trade occurs at the relevant bid or ask quote rather than the bid/ask midpoint. When an option is held for more than one day, we average the total bid/ask cost across days so the daily profits over the holding period are affected equally.[11] For the S&P 500 futures hedges, imposing bid/ask costs is more difficult. Unlike the options data, the futures data contain bid/ask quotes only when a bid (ask) quote is above (below) the most recent transaction price. Nonetheless, procedures such as Smith/Whaley (1994) can be used to estimate the spread from transaction prices. Applying this estimator to our data yields an estimate of just over five cents for the nearby futures contract. This estimate coincides with the minimum tick size and seems reasonable. To impose the bid/ask spread on the futures trades, then, we tack on (or subtract) half of this estimate to each trade price.

Figure 3 illustrates the cumulative profits after imposing these costs. Comparing this figure with Figure 2, we see that the bid/ask spread costs have a large effect. The total profits for all of the trading strategies are either close to zero or negative. In addition, the returns appear to be relatively flat across the sample (except for the hedged volatility spread which trends downward). Table 2 summarizes the daily returns and the returns by subsample. Across the entire sample, only two of the strategies generate positive returns (puts and the volatility spread). Their *t*-statistics (or Sharpe ratios), however, are just 0.38 and 0.44, respectively, and much less than that for S&P 100 index returns during this period (1.35). Examining the subsample results, there are few years when a particular trading strategy earns a larger Sharpe ratio than the S&P 100 portfolio. In fact, for four of the six years, the index portfolio outperforms *all* of the strategies. Based on this evidence, it does not appear that we can earn abnormal returns by trading on the implied volatility's forecast bias.

Though not shown in the table, by far the largest component of the bid/ask costs is for the at-the-money options. These spreads average about 2.5% of option value and range from about 1% to as much as 12%.[12] Our trading rules initiate transactions, on average, about every third day, so this bid/ask cost reduces returns by almost 1% per trading day.[13] Given this information, it might seem that we may be able earn abnormal returns by developing a strategy which requires less trading. But several factors suggest otherwise. First, the transaction costs

[11] This procedure ignores the effect of time value, however, this effect is small because the holding period is typically no more than a few days.
[12] These findings are generally consistent with those reported in George and Longstaff (1993) and Fleming, Ostdiek, and Whaley (1996). George and Longstaff find that the bid/ask spreads for S&P 100 options range from 2% to 20$, and Fleming, Ostdiek, and Whaley find an average spread for at-the-money options of about 4%.
[13] The bid/ask cost is much lower for the out-of-the-money options (used in the hedged volatility spread) and S&P 500 futures (used for delta hedging) which, respectively, reduce returns by about 0.5% and 0.1% per day.

will be relatively high for any strategy because we must pay a bid/ask spread cost (to sell the options) and the maximum holding period is limited (the options expire, on average, in 20 trading days). Second, as we increase the holding period, the variance of returns increases and the number of independent observations decreases, so it becomes more difficult to establish significance. Finally, we need to be sensitive to data snooping concerns.

An alternative strategy might be to focus only on those options that a priori seem most "overpriced." To identify these options, we need to determine when the level of implied volatility bias is greatest. The most straightforward approach is simply to compare the implied volatility to other measures of conditional volatility such as historical volatility or GARCH forecasts. Since these measures contain similar (though not as much) information as the implied volatility, any differences between them may increase with the level of bias. To implement this approach, we estimate the historical volatility and GARCH forecasts for each day in our sample. For the historical volatility, we simply compute the standard deviation of index returns over the previous 28-day period. For the GARCH forecasts, we fit a GARCH(1,1) model across the entire sample, and we use the fitted model to forecast the average volatility over the next 28 days. We use a 28-day period for both estimators to be consistent with the average life of the options used in the trading strategies. Now, for each day, we compute the ratio of the implied volatility to either the historical or GARCH volatility (the "IV ratio"), and we sort our previous trading strategy returns into deciles conditioned on the IV ratio.

Table 3 reports the results. Examining the returns across deciles reveals little evidence that the trading strategy profits covary with the IV ratio. Using the historical volatility, for example, few of the returns are significant and no patterns are apparent across deciles, except perhaps for the call option strategy where the returns for the four largest deciles are negative. Even here, however, the returns are more negative for the fourth decile and the pattern across deciles is sporadic. The results using the GARCH forecasts are similar. Now, the returns tend to be positive in higher deciles for the put option and volatility spread strategies, but again few of the returns are significant and the patterns across deciles are very weak. Thus, we conclude that we cannot earn abnormal returns using this strategy to identify the most biased options.

Finally, we should consider whether the implied volatility bias provides an opportunity for floor traders to earn abnormal profits. Our results indicate that the bias is not large enough to generate profits after transaction costs, but perhaps a floor trader could do so since his transaction costs are relatively low. There are three reasons that this seems unlikely. First, if the trader must initiate trades to exploit the bias, he too must pay the cost of the bid/ask spread. Second, if the trader seeks to profit by simply selling options on customer demand, he still must pay exchange fees, cover his fixed costs, and bear the adverse selection cost of trading with informed traders. These costs may or may not be less than the bid/ask spread, but they certainly reduce the trader's potential profit. Third, as we investigate in the next section, any remaining profits may simply represent compensation for volatility and/or skewness risk. Taken together, these factors may explain why the bias can persist without being competed down to a smaller amount by floor traders.

V. Explanations

Our empirical analysis indicates that S&P 100 option prices, on average, seem to be overstated. This finding (ignoring transaction costs) demonstrates that the implied volatility bias is not entirely due to measurement error and misspecification, but after imposing transaction costs the bias is not large enough to earn abnormal profits. In this section, we examine possible explanations for these results. First, we consider the distribution of the trading strategy profits and evaluate whether the results may be due to sample selection or a risk premium demanded for skewness and higher-order moments of returns; second, we examine whether the profits represent compensation for volatility risk; and, finally, we assess the effect of transaction costs and other market imperfections and the impact of the 1987 stock market crash.

A. Nonnormality of Returns

The distribution of returns is an important issue in evaluating any option trading strategy due to the nature of option payoffs. An option has limited liability and the potential for large profits which indicates that the returns distribution is heavily skewed and leptokurtic. Our strategy of hedging the option positions reduces the asymmetry in returns but does not eliminate it. We can never identify the true valuation model, and, as Merton (1976) argues, if the hedge is based on a misspecified model then the payoff on the hedged position will also be asymmetric.

Table 4 summarizes the distributions of our trading strategy returns in terms of their departures from normality. To read the table, consider the probability (0.0067) reported for the 1% level of S&P 100 returns. This value is determined by finding the return that defines the 1st percentile of the distribution, and then computing the probability of realizing a value this small if returns were normally distributed. Because the probability is less than 1%, it indicates the distribution has a fatter left tail than the normal. The probability for the 99% level (0.9952) exceeds 99%, indicating a fatter right tail than the normal. Therefore, these findings suggest S&P 100 returns exhibit excess kurtosis. For the trading strategies, the left-tail probabilities are even lower than for S&P 100 returns, but the right-tail probabilities are also lower. In addition, the probabilities for the 50% level are all greater than 0.50. These findings indicate that the distribution of trading strategy returns is leptokurtic and negatively skewed. The strategies tend to make money most of the time and occasionally incur large losses.

Next, we consider the significance of the departures from normality. The test statistics, also reported in Table 4, are based on the analysis of variance test developed by Shapiro and Wilk (1965). As the likelihood of normality increases, the Shapiro/Wilk *W*-statistic approaches 1.0. The significance levels depend on sample size and can be computed using the approximation developed by Royston (1982). For ease of interpretation, we also report the corresponding *z*-statistic (which is standard normal) using Royston's transformation. These test statistics indicate strong rejections of normality for all of the return series. The rejection for S&P 100 returns is consistent with the nonnormality of stock returns first reported in Mandelbrot (1963). The rejections for the trading strategy returns are far more severe.

This evidence confirms that the option returns are asymmetric even after hedging. Now, we must address two concerns that this finding motivates regarding the significance of the trading strategy profits. The first is sample selection. Our short positions may appear to earn abnormal returns only because the choice of sample period poorly approximates the true distribution. There are several reasons, however, to believe this explanation is incomplete. First, our sample is relatively large, consisting of six years of daily data.[14] Second, the average daily profits (without transaction costs) are consistently positive year after year, even when the occasional large losses occur. Finally, as we illustrate below, even if we include the 1987 crash in our sample, all of the losses incurred are more than offset during the remainder of the period.

The second concern is that the nonnormality may be related to the economic nature of the profits. Our reliance on the Sharpe ratio in Section IV to evaluate performance assumes that investors care only about mean and variance. But the option returns are negatively skewed

[14] Most prior research is based on one to three year sample periods. See Bates (1996) for a survey of this research and a more thorough discussion of the concerns about sample size in evaluating option trading strategies.

and leptokurtic, and investors may demand a risk premium for exposure to higher-order moments of returns. An investor with power utility, for example, values skewness because the third derivative of his utility function is positive. In this case, the returns on our option trading strategies (which sell skewness) should be higher relative to a world in which only mean and variance matters. Thus, they may offer a high mean/variance combination only because they also entail the adverse possibility of large losses (i.e., the negative skewness).

To assess the merit of this argument, we reexamine our results using Leland's (1996) modified-beta approach for evaluating performance. Leland assumes that the marginal investor has power utility and the market portfolio follows an *iid* binomial process such that in the limit its log price relatives converge to a normal distribution, consistent with Black/Scholes (1973) and Merton (1973). Under these conditions, the expected return, $E(r_p)$, for an asset with any arbitrary distribution of returns must satisfy,

$$E(r_p) = r_f + B_p[E(r_m) - r_f], \tag{7}$$

where r_f is the riskless interest rate, r_m is the return on the market portfolio,

$$B_p = \frac{\text{Cov}[r_p, -(1+r_m)^{-g}]}{\text{Cov}[r_m, -(1+r_m)^{-g}]}, \tag{8}$$

and

$$g = \frac{\ln[E(1+r_m)] - \ln(1+r_f)}{\text{Var}[\ln(1+r_m)]}. \tag{9}$$

We implement this approach using the S&P 100 index to proxy for the market portfolio. First, we estimate the parameters of g using the realized daily r_m over our sample and the daily series of 30-day riskless interest rates as r_f. Then, we use the time-series of $-(1+r_m)^{-g}$ to compute the covariances needed to estimate B_p. For comparison, we also compute the ordinary betas, b_p, for each trading strategy by using the daily r_m instead of $-(1+r_m)^{-g}$ in equation (8).

Table 5 reports the results. Examining the ordinary beta estimates first, we see that the call option strategy is positively correlated with S&P 100 returns, but the other three strategies have negative betas. As a result, the risk-adjusted returns for the call option strategy (reported as a in the table) seem less abnormal (t-statistic = 3.76) than the excess returns reported in Table 1 (t-statistic = 3.99), while the risk-adjusted returns for the other strategies now seem even more significant. The final three columns report the results for the modified beta estimates. As expected, the modified betas for each strategy are greater than the ordinary betas, reflecting the compensation required for negative skewness in the option strategies. However, the largest change occurs for the call option strategy, and this change is just 0.35. Multiplying this amount by the excess return on the market portfolio (0.0313, from Table 1) reveals that using the modified beta has only a trivial effect on our evaluation of the trading strategy returns. The modified alpha (A) for each strategy remains significantly greater than zero.

These results indicate that the risk premium associated with the skewness and higher-order moments of the option strategy returns accounts for just a small amount of the apparent bias. We should note that this analysis requires a number of assumptions and the results may vary with a different utility function or different proxy for the market portfolio. It seems unlikely, however, that any reasonable changes in methodology would qualitatively alter the conclusion. To do this, we must explain why the put option strategy, whose returns are strongly negatively correlated with S&P 100 index returns, earns significantly positive excess returns.

B. Volatility Risk Premium

Another possibility is that the profits represent compensation for volatility risk. S&P 100 returns and volatility are negatively correlated which implies that an “investment” in volatility tends to pay off when the market (a large subset of aggregate wealth) is low and the marginal utility of wealth is high. This investment has a negative “beta” and should earn a negative risk premium. Since our trading strategies sell options (or volatility), a negative volatility risk premium increases the expected returns relative to a world in which volatility risk is not priced. Therefore, the strategies may appear to earn abnormal returns only because we have not accounted for their volatility risk.

One implication of this argument is that the strategies with greater volatility exposure should earn higher profits. To verify this, we rank order the four strategies by their vegas (i.e., ¶f_t/¶s_t) under the Black/Scholes model. The vegas for a call and a put option are the same, but the strategy vegas differ due to the quantities traded. Each strategy requires a $100 investment, and because an at-the-money call has greater value than an at-the-money put, the put option strategy requires more options and has a higher vega. The volatility spread uses both calls and puts, so its vega is between the call and put option strategy vegas. The hedged volatility spread has the lowest vega because this strategy also requires the purchase of out-of-the-money options. Comparing this ordering with the mean returns reported in Table 1 reveals that the strategies with greater (negative) volatility risk indeed earn the greater returns. The put option strategy earns the highest return (1.1687), followed by the volatility spread (1.0240), call option (0.9858), and hedged volatility spread (0.6517) strategies.

To form a more rigorous test, we now assess whether the returns vary with time-variation in volatility exposure and the volatility risk premium. We can show that the vega for an at-the-money index option increases in volatility for moderate levels of volatility. For extremely high levels of volatility, the relation reverses but is relatively flat. This pattern in volatility exposure suggests that the trading strategy returns should be higher when volatility is higher. We might expect a similar effect due to variation in the volatility risk premium. Option pricing models with stochastic volatility typically assume that the risk premium varies with the level of volatility.[15] As Wiggins (1987) demonstrates, this assumption can be easily justified in a Cox/Ingersoll/Ross (1985) framework. If the premium increases (i.e., becomes more negative) when volatility is higher, then the trading strategy returns should be even higher.

To evaluate this hypothesis, we use the regression,

$$r_{i,t} = a_i + b_i \bar{s}_t + e_{i,t}, \tag{10}$$

where $r_{i,t}$ is the day t return for trading strategy i.[16] Unfortunately, the nonnormality of returns suggests that parameter estimates obtained using OLS may be inefficient. There are a large number of "outliers" which receive disproportionate weight in OLS, causing the estimates to be extremely sensitive. We alleviate this concern by using a robust *M*-estimator based on Windsorized residuals, $e^*_{i,t} = e_{i,t}$ for $-cs_i < e_{i,t} < cs_i$, and otherwise $e^*_{i,t} = cs_i \times \text{sign}(e_{i,t})$, where s_i is the standard error and c is an arbitrary constant. The procedure treats small residuals ($|e_{i,t}| < cs_i$) as in OLS but limits the influence of large ones ($|e_{i,t}| > cs_i$). Because the results depend on the choice of c, we consider two reasonable values, $c = 1.50$ and $c = 1.75$. Huber (1981) describes the estimation algorithm and the asymptotic properties of the estimates.[17]

Table 6 summarizes the regression results based on the trading strategy returns before imposing transaction costs. In general, the coefficient estimates show only marginal significance and the regressions have relatively low explanatory power (not shown in the table, but

[15] See, for example, Melino and Turnbull (1990) and Heston (1993).
[16] Our analysis suggests the relation between volatility and returns may be nonlinear. To allow this possibility, we also considered a polynomial series expansion of equation (10). Andrews (1991) demonstrates that under reasonable regularity conditions the expansion can approximate any functional form to an arbitrary degree of accuracy. The results, however, indicate that the higher-order terms contribute little explanatory power.
[17] See Judge, et al. (1985), Section 20.4, for a more general description of the procedure.

less than 1%). Nonetheless, the estimates of *b* are positive for all of the trading strategies, consistent with the hypothesis that returns increase with volatility. The estimates are greater for the first three strategies in the table which also seems reasonable because, as we explained above, the hedged volatility spread entails much less volatility exposure. These results support the idea that the trading strategy returns may be due, in part, to a volatility risk premium that varies with the level of volatility. This is consistent with Lamoureux and Lastrapes (1993) who find a similar result for the implied volatility forecast bias (rather than prices) of stock options.

We ignore transaction costs in the returns used in Table 6 because the bid/ask spread also varies with the level of volatility. Specifically, the spread should increase because market makers are risk averse and the adverse selection cost increases with volatility. This relation has been empirically documented by Harris (1994) and others, and is apparent in our data. Dividing the sample in half based on the level of implied volatility, the average bid/ask spread for nearby, at-the-money options is about 10% higher in the subsample with higher volatility. Whether this variation positively or negatively affects the compensation for volatility risk depends on who initiates the trading. As we argue below, there is a natural demand for buying index options, and trades initiated at the ask price provide even more compensation for volatility risk when the level of volatility is higher.

C. Transaction Costs and Other Market Imperfections

Transaction costs and other market imperfections may also contribute to the observed bias. The trading strategy profits disappear after imposing bid/ask transaction costs, and this suggests the bias may exist precisely because of them. Bid/ask spreads allow option prices (and implied volatilities) to drift from their "true" values without permitting arbitrage opportunities. For example, suppose the "true" expected volatility is 18%. The corresponding bid and ask prices for might imply volatilities of 17.5% and 18.5%, respectively, but the bid (or ask) volatility can drift as high as 17.99 (or as low as 18.01) without allowing arbitrage profits. As Figlewski (1997) explains, this flexibility allows market markers to adjust their quotes in response to supply and demand conditions even though volatility expectations may not have changed.

The flexibility allowed by the bid/ask spread is limited because at-the-money option prices are relatively sensitive to the volatility rate. However, the flexibility increases after we include other transaction costs and market imperfections. Suppose the "true" volatility moves outside the bid/ask band, and that to earn "arbitrage" profits we sell options and delta hedge the position. Delta hedging with the underlying S&P 100 index portfolio is extremely costly and faces substantial execution risk. Alternatively, we can reduce transaction costs by hedging with S&P 500 futures, but now the hedge entails basis and tracking risks. In addition, even if we adequately hedge the position against market risk, we still face the uncertainty that volatility may change. Therefore, the apparent "arbitrage" can be costly to execute and certainly involves risk. This intuition is consistent with Figlewski (1989).

These arguments suggest that S&P 100 option prices can systematically (or randomly) deviate from investor volatility expectations without generating equilibriating trades. This, in turn, would weaken the implied volatility as a forecast. One implication is that the implied volatility's forecast performance should be better in markets where the arbitrage is less difficult to implement. Figlewski (1997) finds some evidence to support this position. A second implication is that the implied volatility's tendency to overstate or understate volatility may be altered when a shock occurs that systematically affects supply and demand.

Anecdotal evidence suggests this type of event occurred after the 1987 stock market crash. The crash revealed imperfections in dynamic portfolio insurance, and increased demand for explicit portfolio insurance. As Dumas, Fleming, and Whaley (1998) argue, this institutional buying pressure may have bid up the price of out-of-the-money puts until market makers were eventually willing to step in and accept the "bet." Given the cross-sectional relation among option prices, this may have increased at-the-money option prices as well. This implication suggests that the profits earned by our trading strategies may have increased after the crash.

The crash's impact on the structure of index option prices has been documented in prior research. Rubinstein (1994) shows that the volatility smile changed shape and has since been more of a skew. Bates (1998) shows that the skewness premium, the deviation between prices for call and put options of the same moneyness, permanently increased. One explanation for these findings is that the crash increased the market's assessment of another "crash." In terms of volatility, this implies that investors systematically anticipated greater volatility than actually occurred. This tendency could persist due to transaction costs and other market imperfections, as we explained above. Another possibility, however, is that the crash changed investors' risk aversion, increasing the premium demanded for equity, volatility, and/or jump risk. Finally, the crash may have revealed information about the number of portfolio insurers, and, as Grossman and Zhou (1996) argue, this could increase the correlation between index returns and volatility.

We can use our option trading strategies to assess the validity of these explanations. Under each explanation, the degree of implied volatility bias may have changed following the crash, but these changes should only be apparent in our trading strategies if they are due to transaction costs or changes in the risk premiums. Under the final explanation, the bias changes only because the degree of model misspecification increases. To evaluate these effects, we use the same trading strategies and methodology as before, but we now extend the sample period back to the beginning of 1986. Then, we assess whether any changes in returns are apparent around the time of the crash.

Figure 4 shows the cumulative profits for each trading strategy from 1986 through 1990. The profit curves are relatively flat for two years leading up to the crash, fluctuating randomly with no clear trend. Then, during the crash, all of the strategies experienced large losses except for the hedged volatility spread. After the crash, each strategy began to earn steady positive profits (in the absence of transaction costs), and these profits continued through the remainder of the sample (see Figure 2). This evidence suggests that the crash indeed had an effect on the implied volatility's forecast bias, just as it did on the volatility skew and the skewness premium. Moreover, because the bias is apparent in our option trading strategies, we conclude that this change occurred due to risk premium effects and/or fears concerning another crash.

V. Conclusions

We set out to explain why the implied volatility for S&P 100 options systematically overstates future stock market volatility. We consider three possible explanations. At one extreme, the bias may be purely the result of systematic measurement error and/or misspecification of the implied volatility as equal to the market's volatility forecast. At the other extreme, the bias may result from inefficiency in the option market. In between these alternatives, the bias may be due to transaction costs and other market imperfections, a risk premium for volatility and/or jump risk, or investor preferences for skewness and higher-order moments of returns. These factors can allow option values to differ from their theoretical values in a frictionless risk-neutral world without signaling the opportunity for abnormal profits.

To distinguish among these possibilities, we examine the profits for trading strategies designed to exploit the implied volatility's observed bias. In the absence of transaction costs, the strategies earn large positive profits with reward/risk ratios three to five times greater than the S&P 100 index. Moreover, the profits are systematic and not isolated to a particular period. This evidence suggests that the implied volatility bias is not purely the result of measurement and misspecification, but that it is truly a function of option market prices. After imposing bid/ask transaction costs, the trading profits disappear. Therefore, the bias does not seem large enough to signal option market inefficiency.

We next examine the nature of the trading profits to evaluate possible explanations for these results. First, the distribution of returns is asymmetric, so the profits may represent compensation for skewness and higher-order moments of returns, and they may be sensitive to the choice of sample period. We find, however, that these risks explain only a small part of the trading strategy returns and that the returns are still significant if we include the 1987 crash. Second, the observed bias may be caused by a negative risk premium for volatility and/or jump risk. Such a premium seems reasonable, and the trading profits increase with greater volatility exposure and higher volatility. Third, it is difficult to exploit volatility "mispricing" due to transaction costs and other market imperfections. These frictions allow option prices to systematically deviate from their true values without signaling arbitrage opportunities. Finally, the trading strategy profits do not appear until after the 1987 crash, so the bias seems related to the structural change observed in index option prices following this event.[18] Such a change is consistent with both the risk premium and trading cost explanations above.

These findings have two main implications for current research. First, there is growing evidence that the implied volatility may be useful in forecasting stock market volatility. But, as Figlewski (1997) describes, we must account for the bias to maximize the implied volatility's forecasting value. Our results reveal some likely sources to pursue in seeking a bias correction. Second, our evidence regarding a volatility risk premium substantiates the importance of this parameter for option valuation. Current research has well-documented the deficiencies of the Black/Scholes model, and a number of more complex models have been developed as alternatives. These models, however, are often less reliable in practice because they involve more free parameters that must be estimated. Our results indicate that if there is a priority for including these additional parameters, the volatility risk premium should receive added emphasis.[19]

[18] See Rubinstein (1994) and Bates (1998).
[19] This conclusion is consistent with Bates (1998) who finds that a volatility risk premium accounts for the difference between realized volatility and the volatility expectation implied by option prices.

References

Andrews, D.W.K., 1991, Asymptotic normality of series estimators for nonparametric and semi-nonparametric regression, *Econometrica* 59, 307–346.

Bates, D.S., 1996, Testing option pricing models, in G.S. Maddala and C.R. Rao, ed., *Handbook of Statistics: Statistical Methods in Finance* 14, Amsterdam: Elsevier, 567–611.

Bates, D.S., 1998, Post-‘87 crash fears in the S&P 500 futures options market, *Journal of Econometrics*, Forthcoming.

Black, F., 1976, Studies of stock price volatility changes, *Proceedings of the 1976 Meetings of the American Statistical Association*, Business and Economic Statistics Section, 177–181.

Black, F. and M.S. Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–659.

Canina, L. and S. Figlewski, 1993, The information content of implied volatility, *Review of Financial Studies* 6, 659–681.

Cox, J., J. Ingersoll, and S. Ross, 1985, An intertemporal general equilibrium model of asset prices, *Econometrica* 53, 363–384.

Day, T.E. and C.M. Lewis, 1988, The behavior of the volatility implicit in the prices of stock index options, *Journal of Financial Economics* 22, 103–122.

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

Dumas, B., J. Fleming, and R.E. Whaley, 1998, Implied volatility functions: Empirical tests, *Journal of Finance*, Forthcoming.

Feinstein, S., 1989, A theoretical and empirical investigation of the Black-Scholes implied volatility, Ph.D. dissertation (Yale University, New Haven, CT).

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

Figlewski, S., 1997, Forecasting volatility, *Financial Markets, Institutions and Instruments* Vol. 6.

Fleming, J., 1998, The quality of market volatility forecasts implied by S&P 100 index option prices, *Journal of Empirical Finance*, Forthcoming.

Fleming, J., B. Ostdiek, and R.E. Whaley, 1995, Predicting stock market volatility: A new measure, *Journal of Futures Markets* 15, 265–302.

Fleming, J., B. Ostdiek, and R.E. Whaley, 1996, Trading costs and the relative rates of price discovery in stock, futures, and option markets, *Journal of Futures Markets* 16, 353–387.

Fleming, J. and R.E. Whaley, 1994, The value of wildcard options, *Journal of Finance* 49, 215–236.

French, K.R., G.W. Schwert, and R.F. Stambaugh, 1987, Expected stock returns and volatility, *Journal of Financial Economics* 19, 3–29.

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

Grossman, S.J., and Z. Zhou, 1996, Equilibrium analysis of portfolio insurance, *Journal of Finance* 51, 1379–1403.

Harris, L.E., 1994, Minimum price variations, discrete bid-ask spreads, and quotation sizes, *Review of Financial Studies* 7, 149–178.

Harvey, C.R. and R.E. Whaley, 1992a, Market volatility prediction and the efficiency of the S&P 100 index option market, *Journal of Financial Economics* 31, 43–73.

Harvey, C.R. and R.E. Whaley, 1992b, Dividends and S&P 100 index option valuation, *Journal of Futures Markets* 12, 123–137.

Heston, S.L., 1993, A closed-form solution for options with stochastic volatility with applications to bond and currency options, *Review of Financial Studies* 6, 327–343.

Huber, P.J., 1981, Robust statistics (Wiley: New York).

Hull, J. and A. White, 1987, The pricing of options on assets with stochastic volatilities, *Journal of Finance* 42, 281–300.

Jackwerth, J., and M. Rubinstein, 1996, Recovering probability distributions from option prices, *Journal of Finance* 51, 1611–1631.

Jorion, P., 1995, Predicting volatility in the foreign exchange market, *Journal of Finance* 50, 507–528.

Judge, G.G., W.E. Griffiths, R.C. Hill, H. Lütkepohl, and T-C. Lee, 1985, The theory and practice of econometrics, 2nd Ed (Wiley, New York).

Lamoureux, C.G. and W.D. Lastrapes, 1993, Forecasting stock return variance: Toward an understanding of stochastic implied volatilities, *Review of Financial Studies* 6, 293–326.

Leland, H.E., 1996, Beyond mean-variance: Performance measurement of portfolios using options or dynamic strategies, Working paper, University of California at Berkeley.

Mandelbrot, B., 1963, The variation of certain speculative prices, *Journal of Business* 36, 394–419.

Melino, A., and S.M. Turnbull, 1990, Pricing foreign currency options with stochastic volatility, *Journal of Econometrics* 45, 239–265.

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

Merton, R.C., 1976, Option pricing when underlying stock returns are discontinuous, *Journal of Financial Economics* 3, 125–144.

Royston, J.P., 1982, An extension of Shapiro and Wilk's *W* test for normality to large samples, *Applied Statistics* 31, 115–124.

Rubinstein, M., 1994, Implied binomial trees, *Journal of Finance* 49, 771–818.

Scott, L.O., 1987, Option pricing when the variance changes randomly: Theory, estimation, and an application, *Journal of Financial and Quantitative Analysis* 22, 419–438.

Shapiro, S.S. and M.B. Wilk, 1965, An analysis of variance test for normality (complete samples), *Biometrika* 52, 591–611.

Sharpe, William F., 1966, Mutual fund performance, *Journal of Business* 39, 119–138.

Sharpe, William F., 1994, The Sharpe ratio, *Journal of Portfolio Management* 21, 49–58.

Smith, T. and R.E. Whaley, 1994, Estimating the effective bid-ask spread from time and sales data, *Journal of Futures Markets* 4, 437–455.

Whaley, R.E., 1986, Valuation of American futures options: Theory and empirical tests, *Journal of Finance* 41, 127–150.

Wiggins, J., 1987, Options values under stochastic volatility: Theory and empirical estimates, *Journal of Financial Economics* 19, 351–372.

Table 1
Daily Profits for Short S&P 100 Option Trading Strategies Ignoring Transaction Costs

The table reports the average daily profits for four short S&P 100 option trading strategies in the absence of transaction costs. The options sold are the nearby (but at least 15 days to expiration), at-the-money call and put options, and we assume that each day's trades occur at the average of the bid/ask quotes available at 3:00 CST. The first two strategies simply sell the call or put option, respectively, and then delta hedge the option using S&P 500 futures. The third strategy is a short volatility spread (i.e., sell both call and put options), and the fourth strategy is a volatility spread hedged by buying deep out-of-the-money (30 index points) call and put options. For comparison, we also consider the profits for investing in the S&P 100 index. The profits reported in the table are based on constant $100 daily investments, and are expressed net of the 30-day riskless interest rate. Below the average profits for each strategy, we report the sample standard deviation of the daily profits and the *t*-statistic for the difference of the average profit from zero. The sample period is January 4, 1988 through December 31, 1993.

	Entire Sample	By Year					
		1988	1989	1990	1991	1992	1993
Observations	1515	251	252	253	252	254	253
S&P 100 Index							
Mean Return	0.0313	0.0210	0.0736	–0.0325	0.0824	0.0117	0.0322
Standard Deviation	0.0232	0.0734	0.0556	0.0669	0.0601	0.0399	0.0351
t-statistic	1.35	0.29	1.32	–0.49	1.37	0.29	0.92
Call Options							
Mean Return	0.9858	0.6778	1.2750	0.8049	0.6942	1.6721	0.7857
Standard Deviation	0.2471	0.6819	0.8483	0.5036	0.4736	0.4134	0.6116
t-statistic	3.99	0.99	1.50	1.60	1.47	4.04	1.28
Put Options							
Mean Return	1.1687	1.0730	0.7668	1.4437	0.5969	1.7444	1.3805
Standard Deviation	0.2531	0.5088	1.0478	0.4219	0.4854	0.4431	0.5900
t-statistic	4.62	2.11	0.73	3.42	1.23	3.94	2.34
Volatility Spread							
Mean Return	1.0240	0.9473	0.9468	1.2350	0.3945	1.6916	0.9227
Standard Deviation	0.1546	0.4257	0.3968	0.3696	0.4032	0.2554	0.3976
t-statistic	6.62	2.23	2.39	3.34	0.98	6.62	2.32
Hedged Volatility Spread							
Mean Return	0.6517	0.3553	0.6978	0.9578	0.0363	1.2471	0.6089
Standard Deviation	0.1458	0.4002	0.3717	0.3611	0.3512	0.2577	0.3837
t-statistic	4.47	0.89	1.88	2.65	0.10	4.84	1.59

Table 2
Daily Profits for Short S&P 100 Option Trading Strategies after Bid/Ask Transaction Costs

The table reports the average daily profits for four short S&P 100 option trading strategies after imposing bid/ask transaction costs. The trading strategies are: (1) short call and (2) short put options delta hedged using S&P 500 futures, (3) short volatility spread, and (4) hedged volatility spread. For each strategy, the options sold are the nearby (but at least 15 days to expiration), at-the-money call and/or put options, and we impose bid/ask transaction costs by assuming that each day's trades occur at the relevant bid or ask quote available at 3:00 CST, and that each day's futures trades (used in delta hedging) cost $0.05 per contract. The profits reported in the table are based on constant $100 daily investments, and are expressed net of the 30-day riskless interest rate. Below the average profits for each strategy, we report the sample standard deviation of the daily profits and the *t*-statistic for the difference of the average profit from zero. The sample period is January 4, 1988 through December 31, 1993.

	Entire Sample	By Year					
		1988	1989	1990	1991	1992	1993
Observations	1515	251	252	253	252	254	253
S&P 100 Index							
Mean Return	0.0313	0.0210	0.0736	–0.0325	0.0824	0.0117	0.0322
Standard Deviation	0.0232	0.0734	0.0556	0.0669	0.0601	0.0399	0.0351
t-statistic	1.35	0.29	1.32	–0.49	1.37	0.29	0.92
Call Options							
Mean Return	–0.0682	–0.4091	0.2894	–0.2894	–0.4973	0.6341	–0.1428
Standard Deviation	0.2551	0.7207	0.8657	0.5188	0.4892	0.4305	0.6246
t-statistic	–0.27	–0.57	0.33	–0.56	–1.02	1.47	–0.23
Put Options							
Mean Return	0.0982	0.0366	–0.2785	0.3307	–0.6458	0.7040	0.4350
Standard Deviation	0.2595	0.5246	1.0688	0.4353	0.4963	0.4563	0.6085
t-statistic	0.38	0.07	–0.26	0.76	–1.30	1.54	0.71
Volatility Spread							
Mean Return	0.0707	–0.0073	0.0443	0.2332	–0.7062	0.7549	0.0989
Standard Deviation	0.1623	0.4503	0.4131	0.3889	0.4207	0.2748	0.4143
t-statistic	0.44	–0.02	0.11	0.60	–1.68	2.75	0.24
Hedged Volatility Spread							
Mean Return	–0.9803	–1.2781	–0.8426	–0.8161	–1.8506	–0.3269	–0.7755
Standard Deviation	0.1605	0.4495	0.4023	0.3932	0.3879	0.2941	0.4133
t-statistic	–6.11	–2.84	–2.09	–2.08	–4.77	–1.11	–1.88

Table 3
Short S&P 100 Option Trading Profits Conditioned on the Level of Implied Volatility Relative to Historical Volatility and GARCH Volatility Forecasts

The table reports the average daily profits for the four short S&P 100 option trading strategies after conditioning on the ratio of the daily implied volatility to either (a) the 28-day historical volatility, or (b) the forecasted 28-day average volatility under a GARCH(1,1) model. The trading strategies are: (1) short call and (2) short put options delta hedged using S&P 500 futures, (3) short volatility spread, and (4) hedged volatility spread. For each strategy, the options sold are the nearby (but at least 15 days to expiration), at-the-money call and/or put options, and we impose bid/ask transaction costs by assuming that each day's trades occur at the relevant bid or ask quote available at 3:00 CST. The profits reported for each strategy are based on constant $100 daily investments, and are expressed net of the 30-day riskless interest rate. Next to the average profits for each strategy, we report the *t*-statistic for the difference of the average profit from zero. The sample period is January 4, 1988 through December 31, 1993.

		S&P 100		Call Options		Put Options		Vol. Spread		Hedged Spread	
Decile	Decile Range	Mean	*t*-stat	Mean	*t*-stat	Mean	*t*-stat	Mean	*t*-stat	Mean	*t*-stat
Implied Volatility / Historical Volatility Deciles											
1	0.61 – 0.92	0.01	0.10	0.09	0.12	–1.09	–0.72	–0.19	–0.30	–1.12	–1.80
2	0.92 – 1.01	0.12	1.70	0.19	0.32	0.48	0.89	0.29	0.68	–1.23	–2.70
3	1.01 – 1.09	–0.03	–0.44	0.64	0.83	0.11	0.17	0.09	0.18	–0.93	–1.94
4	1.09 – 1.14	–0.06	–0.73	–0.69	–0.94	–1.38	–1.93	–1.18	–1.89	–1.85	–3.56
5	1.14 – 1.19	0.01	0.23	0.38	0.68	1.17	2.46	0.95	2.78	–0.34	–0.94
6	1.19 – 1.27	0.04	0.63	0.16	0.27	–0.42	–0.57	0.28	0.66	–0.62	–1.54
7	1.27 – 1.35	0.02	0.27	–0.50	–0.68	–0.01	–0.02	–0.06	–0.12	–0.92	–1.97
8	1.35 – 1.43	0.16	2.20	–0.04	–0.07	–0.32	–0.48	0.05	0.10	–1.14	–1.87
9	1.43 – 1.56	0.02	0.31	–0.48	–0.43	0.65	1.07	0.19	0.37	–0.80	–1.41
10	1.56 – 2.46	0.01	0.18	–0.41	–0.33	1.80	1.61	0.26	0.51	–0.85	–1.60
Implied Volatility / GARCH Forecast Deciles											
1	0.65 – 0.83	0.00	0.02	–0.20	–0.23	–1.74	–1.10	–0.11	–0.18	–0.78	–1.29
2	0.83 – 0.90	–0.01	–0.16	0.39	0.55	0.44	0.72	0.29	0.78	–0.66	–1.69
3	0.90 – 0.97	–0.01	–0.10	–1.20	–1.46	–0.56	–0.75	–0.79	–1.37	–1.58	–3.32
4	0.97 – 1.02	0.07	1.18	1.08	1.58	–0.42	–0.70	–0.02	–0.05	–1.04	–2.06
5	1.02 – 1.06	0.03	0.56	0.03	0.05	–0.16	–0.28	–0.16	–0.36	–1.11	–2.63
6	1.06 – 1.10	0.11	1.62	0.69	1.27	0.15	0.28	0.47	1.15	–0.64	–1.50
7	1.10 – 1.15	–0.01	–0.17	–0.95	–0.76	1.72	1.57	0.43	0.83	–0.66	–1.29
8	1.15 – 1.24	0.05	0.76	0.28	0.51	0.41	0.73	0.14	0.28	–1.00	–2.18
9	1.24 – 1.37	0.13	1.41	–0.03	–0.05	0.84	1.23	0.15	0.25	–1.29	–2.19
10	1.37 – 2.10	–0.05	–0.48	–0.77	–0.72	0.29	0.46	0.30	0.52	–1.06	–1.64

Table 4
Distribution of Daily Profits for Short S&P 100 Option Trading Strategies

The table summarizes the distribution of daily trading profits for the four short S&P 100 option trading strategies described in the paper. These strategies are: (1) short call and (2) short put options delta hedged using S&P 500 futures, (3) short volatility spread, and (4) hedged volatility spread. For each strategy, the options sold are the nearby (but at least 15 days to expiration), at-the-money call and/or put options, and we assume that each day's trades occur at either the average of the bid/ask prices (ignoring transaction costs) or at the relevant bid or ask price (including bid/ask transaction costs) available at 3:00 CST. The profits for each strategy are based on constant $100 daily investments, expressed net of the 30-day riskless interest rate. The table is constructed from the distribution of these daily profits over the period from January 4, 1988 through December 31, 1993. The probability values reported for the *i*th percentile represent the probabilities of realizing values as large as those observed in the empirical distribution under the assumption that the daily profits are normally distributed. The table also reports the results of Shapiro-Wilk (1965) tests for normality. The *W*-statistic approaches 1.0 as the likelihood of normality increases, and the *z*-statistic, computed using Royston's (1982) transformation, is distributed standard normal under the null of normality.

		Ignoring Transaction Costs				Including Bid/Ask Transaction Costs			
	S&P 100 Index	Call Options	Put Options	Volatility Spread	Hedged Spread	Call Options	Put Options	Volatility Spread	Hedged Spread
Probability Values									
Minimum	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
1%	0.0067	0.0011	0.0031	0.0000	0.0001	0.0011	0.0040	0.0000	0.0001
5%	0.0639	0.0994	0.0953	0.0472	0.0476	0.0908	0.0975	0.0448	0.0382
10%	0.1241	0.2001	0.2102	0.1546	0.1706	0.1921	0.2174	0.1573	0.1451
50%	0.5005	0.5403	0.5254	0.5681	0.5554	0.5393	0.5301	0.5667	0.5638
90%	0.8781	0.8155	0.7885	0.8166	0.8273	0.8106	0.7867	0.8079	0.8196
95%	0.9421	0.8629	0.8542	0.8718	0.8861	0.8577	0.8460	0.8682	0.8793
99%	0.9952	0.9435	0.9528	0.9561	0.9605	0.9457	0.9379	0.9473	0.9511
Maximum	1.0000	1.0000	1.0000	0.9976	0.9999	1.0000	1.0000	0.9945	0.9987
Shapiro-Wilk Tests									
W-statistic	0.9582	0.7937	0.7331	0.8098	0.8458	0.7625	0.7330	0.8116	0.8453
z-statistic	42.03	551.38	653.77	406.48	302.08	555.30	654.11	401.07	303.48
p-value	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000

Table 5
Daily Risk-Adjusted Returns for the S&P 100 Option Trading Strategies

The table reports the estimated risk and risk-adjusted returns for each of the four option trading strategies. We consider two different measures of risk: the ordinary CAPM beta (*b*), and Leland's (1996) modified beta (*B*) which incorporates the risk and return associated with skewness and higher-order moments of returns. Both betas are estimated using the S&P 100 index to proxy for the market portfolio, and we use the yield on the 30-day T-bill to proxy for the riskless interest rate. The table also reports the alphas, *a* and *A*, corresponding to the two risk measures, respectively, which are computed by subtracting the expected returns given our beta estimates from the realized daily returns for each strategy, and then averaging them over the entire sample period. The reported *t*-statistics measure the significance of the alpha estimates under the null that they equal zero. The sample period is from January 4, 1988 to December 31, 1993.

		Ordinary Beta (*b*)			Modified Beta (*B*)		
Trading Strategy	Excess Return	*b*	*a*	*t*-stat	*B*	*A*	*t*-stat
Ignoring Transaction Costs							
Call Options	0.9858	4.16	0.8553	3.76	4.51	0.8444	3.71
Put Options	1.1687	–2.48	1.2464	5.06	–2.21	1.2381	5.02
Volatility Spread	1.0240	–0.02	1.0247	6.63	0.24	1.0165	6.57
Hedged Volatility Spread	0.6517	–1.03	0.6839	4.75	–0.86	0.6786	4.71
After Bid/Ask Transaction Costs							
Call Options	–0.0682	4.12	–0.1972	–0.83	4.48	–0.2086	–0.88
Put Options	0.0982	–2.46	0.1753	0.69	–2.18	0.1665	0.66
Volatility Spread	0.0707	–0.04	0.0720	0.44	0.23	0.0634	0.39
Hedged Volatility Spread	–0.9803	–1.08	–0.9465	–5.97	–0.88	–0.9527	–6.01

Table 6
Variation in S&P 100 Option Trading Strategy Profits with the Level of Volatility

The table reports the estimation results for the regression,

$$r_{i,t} = a_i + b_i \bar{s}_t + e_{i,t},$$

where $r_{i,t}$ is the day t return for trading strategy i, and $\bar{s}_t$ is the daily implied volatility. The trading strategies are: (1) short call and (2) short put options delta hedged using S&P 500 futures, (3) short volatility spread, and (4) hedged volatility spread. For each strategy, the options sold are the nearby (but at least 15 days to expiration), at-the-money call and/or put options, and we assume that each day's trades occur at the average of the bid/ask quotes available at 3:00 CST (i.e., no transaction costs). The regression estimates are obtained using the Windsorized residuals, $e^*_{i,t} = e_{i,t}$ for $-c\,s_i < e_{i,t} < c\,s_i$, and otherwise $e^*_{i,t} = c\,s_i \times \text{sign}(e_{i,t})$, where s_i is the standard error and c is an arbitrary constant. The estimation algorithm and the asymptotic properties of the estimates are described in Huber (1981). In the table, "Proportion" indicates the proportion of observations in the sample where $|e_{i,t}| \le c\,s_i$. The sample period is January 4, 1988 through December 31, 1993.

Trading Strategy	Proportion	*a*	*t*-stat	*b*	*t*-stat
***c* = 1.50**					
Call Options	0.8528	0.5862	0.94	6.7775	1.77
Put Options	0.8389	0.7778	1.38	5.4302	1.62
Volatility Spread	0.8198	0.9257	2.38	4.8226	2.05
Hedged Volatility Spread	0.8290	1.0193	2.67	1.3001	0.56
***c* = 1.75**					
Call Options	0.8964	0.5455	0.88	6.5985	1.74
Put Options	0.8838	0.6816	1.21	5.6343	1.67
Volatility Spread	0.8838	0.8288	2.11	4.8255	2.04
Hedged Volatility Spread	0.8832	0.9346	2.44	1.4464	0.63

Figure 1
The Forecast Bias of S&P 100 Implied Volatility

The figure illustrates the time-series and distribution of the daily S&P 100 implied volatility forecast errors. The daily implied volatility is computed from the nearby (but at least 15 days to expiration), at-the-money call option prices in a 10-minute window centered around 3:00 CST. The daily forecast error is the difference between the implied volatility and the standard deviation of realized daily S&P 100 returns during the remaining life of the option. Both volatilities are annualized on the basis of 252 trading days per year. Panel A shows the time-series of daily forecast errors, and Panel B shows the frequency distribution. The sample period is January 3, 1986 through December 31, 1993.

Panel A: Time-Series of Implied Volatility Forecast Errors

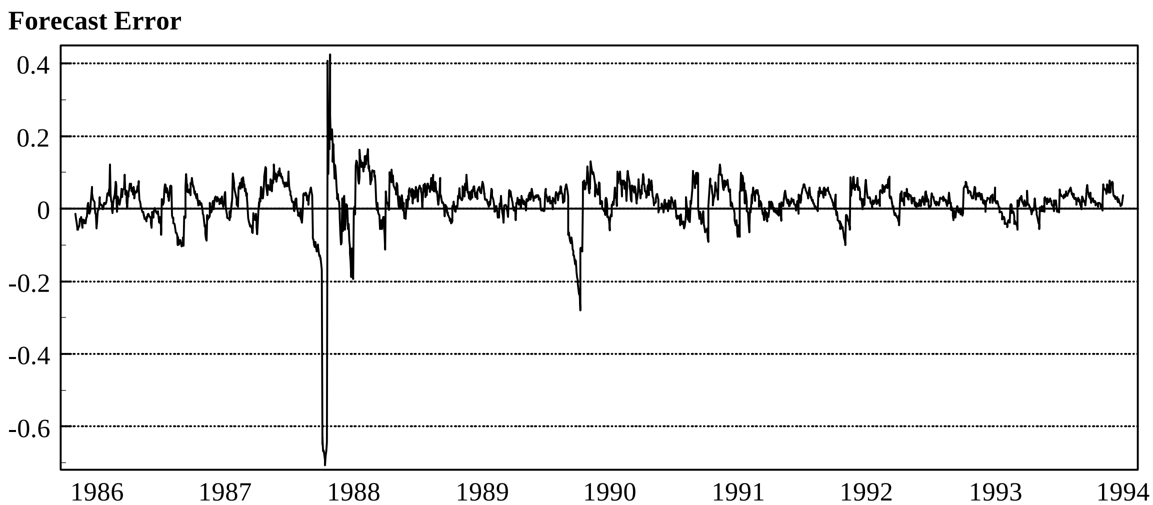

Panel B: Distribution of Implied Volatility Forecast Errors

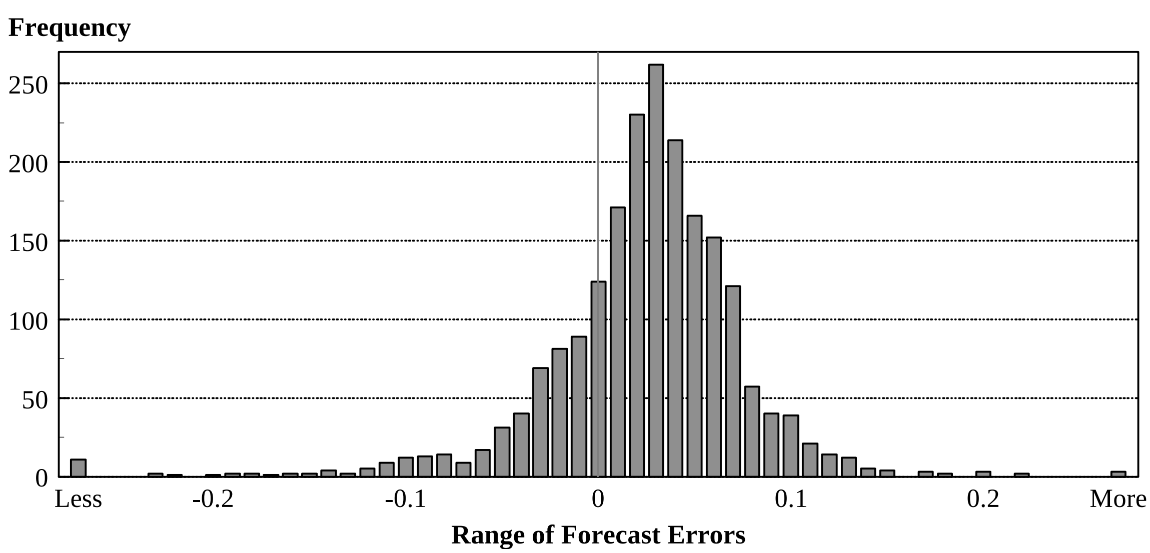

Figure 2
Cumulative Profits for S&P 100 Option Trading Strategies Ignoring Transaction Costs

The figure illustrates the cumulative daily trading profits for the four option trading strategies described in the paper. Each strategy calls for a short position in nearby (but at least 15 days to expiration), at-the-money call and/or put options. Panel A shows the results for the short call and short put option strategies which are delta-hedged using the nearby S&P 500 futures contract. Panel B shows the results for the hedged and unhedged short volatility spread strategies. The hedged volatility spread involves the sale of deep out-of-the-money (30 index points) calls and puts. The daily profits are based on constant $100 investment amounts. The sample period is January 4, 1988 through December 31, 1993.

Panel A: Delta-Hedged Short Option Positions

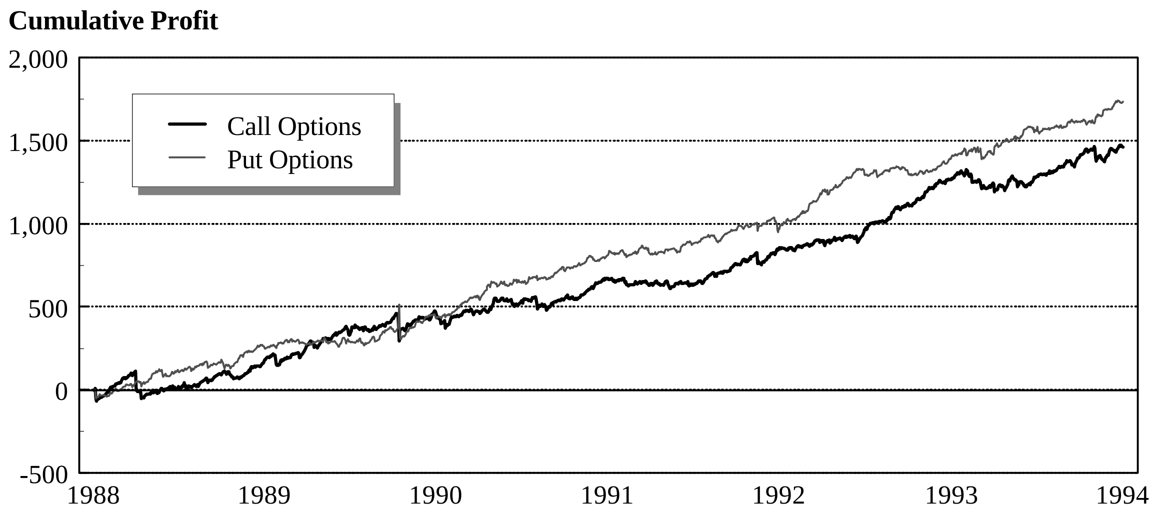

Panel B: Short Volatility Spreads

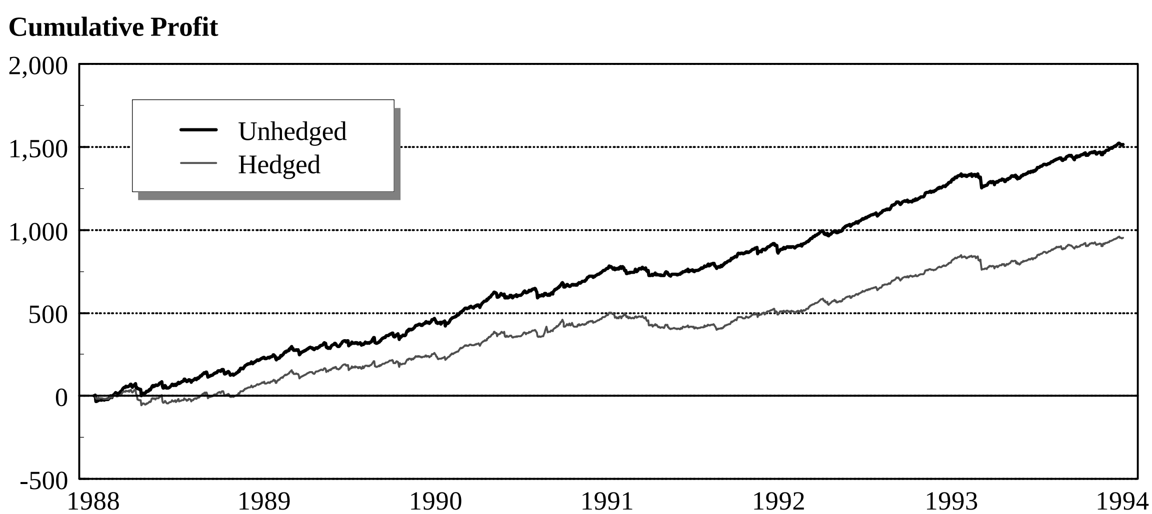

Figure 3
Cumulative Profits for S&P 100 Option Trading Strategies with Bid/Ask Transaction Costs

The figure illustrates the cumulative daily trading profits after bid/ask transaction costs for the four option trading strategies described in the paper. To impose these costs, we assume each option trade occurs at the current bid or ask price and we assume that each futures trade (used in delta hedging) costs $0.05 per contract. Panel A shows the results for the short call and put option strategies delta-hedged using the nearby S&P 500 futures contract. Panel B shows the results for the hedged and unhedged short volatility spread strategies. The daily profits are based on constant $100 investment amounts. The sample period is January 4, 1988 through December 31, 1993.

Panel A: Delta-Hedged Short Option Positions

Cumulative Profit

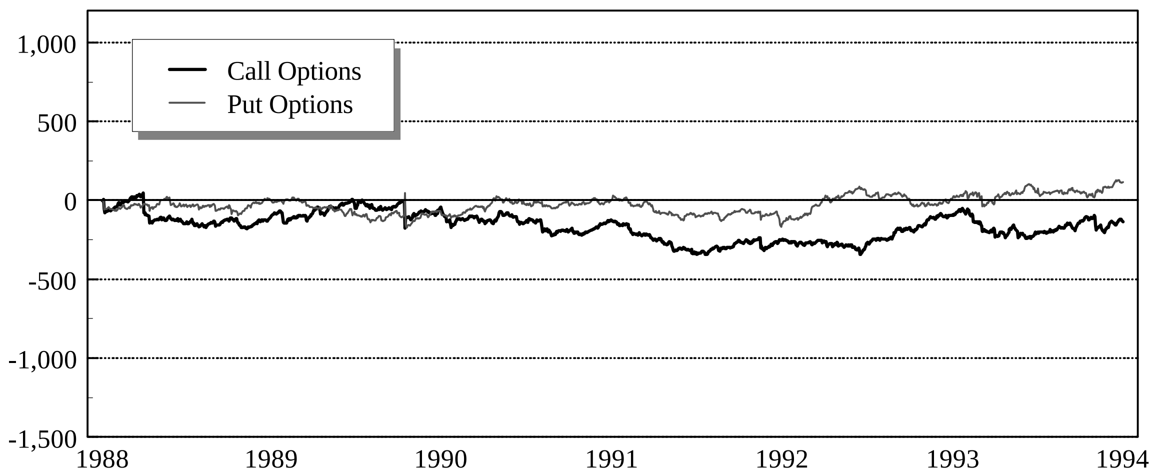

Panel B: Short Volatility Spreads

Cumulative Profit

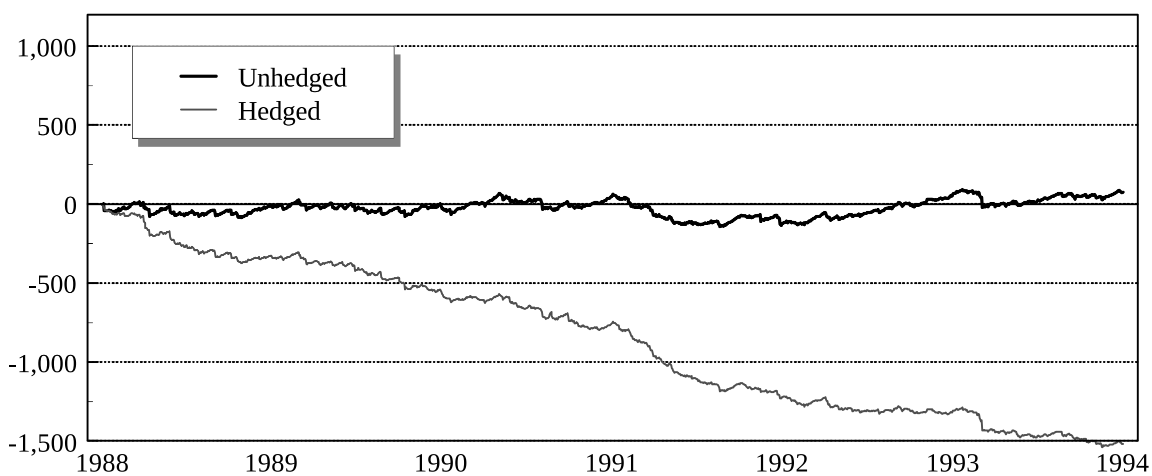

Figure 4
Cumulative Profits for S&P 100 Option Trading Strategies around the 1987 Crash

The figure illustrates the cumulative profits for the four option trading strategies during the period around the 1987 stock market crash. Panel A shows the profits for the short call and put option strategies delta-hedged using the nearby S&P 500 futures contract, and Panel B shows the profits for the hedged and unhedged short volatility spread strategies. The daily profits are based on constant $100 investment amounts and do not include transaction costs. The sample period is January 3, 1986 through December 31, 1991.

Panel A: Delta-Hedged Short Option Positions

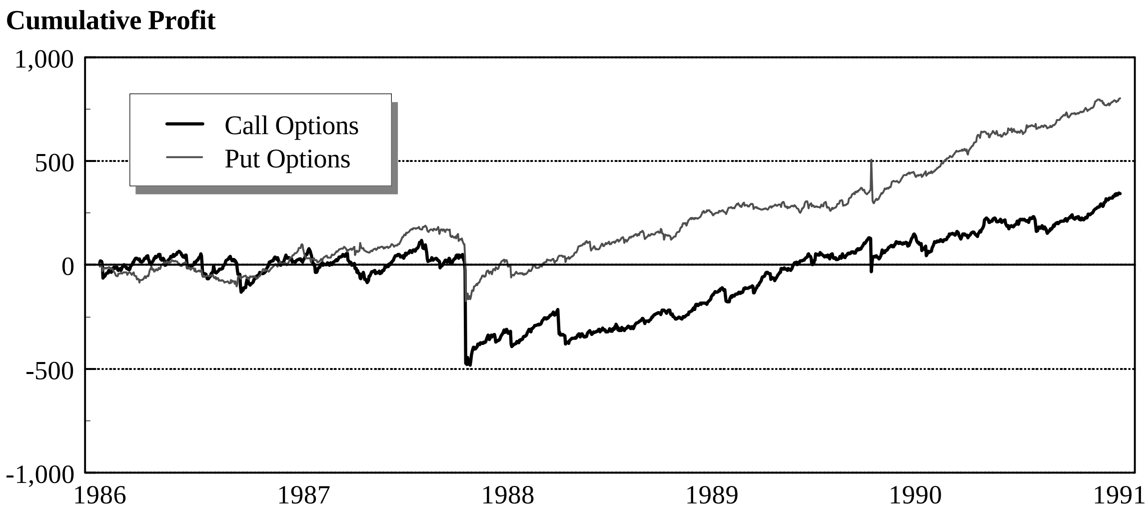

Panel B: Short Volatility Spreads