



American Finance Association

Expected Option Returns

Author(s): Joshua D. Coval and Tyler Shumway

Source: *The Journal of Finance*, Vol. 56, No. 3 (Jun., 2001), pp. 983-1009

Published by: Blackwell Publishing for the American Finance Association

Stable URL: <http://www.jstor.org/stable/222539>

Accessed: 24/11/2009 12:04

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=black>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Blackwell Publishing and American Finance Association are collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Finance*.

<http://www.jstor.org>

Expected Option Returns

JOSHUA D. COVAL and TYLER SHUMWAY*

ABSTRACT

This paper examines expected option returns in the context of mainstream asset-pricing theory. Under mild assumptions, expected call returns exceed those of the underlying security and increase with the strike price. Likewise, expected put returns are below the risk-free rate and increase with the strike price. S&P index option returns consistently exhibit these characteristics. Under stronger assumptions, expected option returns vary linearly with option betas. However, zero-beta, at-the-money straddle positions produce average losses of approximately three percent per week. This suggests that some additional factor, such as systematic stochastic volatility, is priced in option returns.

ASSET-PRICING THEORY CLAIMS that options, like all other risky securities in an economy, compensate their holders with expected returns that are in accordance with the systematic risks they require their holders to bear. Options which deliver payoffs in bad states of the world will earn lower returns than those that deliver their payoffs in good states. The enormous popularity of option contracts has arisen, in part, because options allow investors to precisely tailor their risks to their preferences. With this in mind, a study of option returns would appear to offer a unique opportunity in which to investigate what kinds of risks are priced in an economy. However, although researchers have paid substantial attention to the pricing of options conditional on the prices of their underlying securities, relatively little work has focused on understanding the nature of option returns.

Understanding option returns is important because options have remarkable risk-return characteristics. Option risk can be thought of as consisting of two separable components. The first component is a leverage effect. Because an option allows investors to assume much of the risk of the option's underlying asset with a relatively small investment, options have characteristics similar to levered positions in the underlying asset. The Black-Scholes model implies that this implicit leverage, which is reflected in option betas, should be priced. We show that this leverage should be priced under

* Both authors are from the University of Michigan Business School, Ann Arbor, Michigan. We thank Sugato Bhattacharya, Jonathan Paul Carmel, Artyom Durnev, Nejat Seyhun, René Stulz, an anonymous referee, and seminar participants at the University of Michigan, Michigan State University, Ohio State University, the University of Texas at Austin, and Washington University in Saint Louis for useful comments. We also thank the University of Michigan Business School for purchasing the data used in this study. Remaining errors are, of course, our responsibility.

much more general conditions than the Black–Scholes assumptions. In particular, call options written on securities with expected returns above the risk-free rate should earn expected returns that exceed those of the underlying security. Put options should earn expected returns below that of the underlying security. We present strong empirical support for the pricing of the leverage effect.

The second component of option risk is related to the curvature of option payoffs. Because option values are nonlinear functions of underlying asset values, option returns are sensitive to the higher moments of the underlying asset's returns. The Black–Scholes model assumes that any risks associated with higher moments of the underlying assets returns are not priced because asset returns follow geometric Brownian motion, a two-parameter process. This is equivalent to assuming that options are redundant assets. One testable implication of the redundancy assumption is that, net of the leverage effect, options portfolios should earn no risk premium. In other words, options positions that are delta neutral (have a market beta of zero) should earn the risk-free rate on average. We test this hypothesis, and find quite robust evidence to reject it.

The extraordinary nature of option payoffs has prompted a number of researchers to examine general asset-pricing questions with options data. In fact, some of our conclusions are similar to those of Jackwerth and Rubinstein (1996) and Jackwerth (2000), who use prices to characterize the shape of the risk-neutral density,¹ Canina and Figlewski (1993) and Christensen and Prabhala (1998), who compare historical volatilities to those implied by Black–Scholes, and Merton, Scholes, and Gladstein (1978, 1982), who simulate the returns to a number of options strategies. However, this paper is the first to focus on the theoretical and empirical nature of option returns, rather than volatilities or implied distributions, in the context of broader asset-pricing theory.

We conduct our tests with option returns for three reasons. First, tests based on option returns permit us to directly examine whether the leverage effect is priced without imposing any particular model. Second, analyzing options in risk-return space allows any violations of market efficiency to be directly measured in economic terms. For if options appear mispriced, the economic significance of the mispricing can only be assessed by evaluating the extent to which potential profits outweigh the risks and transaction costs required to exploit the mispricing. Third, by examining option returns, we can consider the economic payoff for taking very particular risks. For example, we can estimate the expected return of a position that is essentially immune to both small market movements and the risk of large market crashes, but is very sensitive to changes in volatility. Tests that measure the impact of particular sources of risk, give a measure of the economic significance of departures from theory, and are somewhat model independent are difficult to construct in any context other than risk-return space.

¹ Jackwerth (2000) analyzes a simulated trading strategy that suggests that put returns may be too low.

Our paper focuses on the returns of call, put, and straddle positions both theoretically and empirically. As stated previously, under very general conditions, calls have positive expected returns and puts have negative expected returns. Moreover, the expected returns of both calls and puts are increasing in the strike price. To test whether option returns are consistent with these implications of the leverage effect, we examine payoffs associated with options on the S&P 500 and S&P 100 indices. We find that call options earn returns that are significantly larger than those of the underlying index, averaging roughly two percent per week. On the other hand, put options receive returns that are significantly below the risk-free rate, losing between eight and nine percent per week. We also find, consistent with the theory, that returns to both calls and puts are increasing in the option strike price.

Although option returns conform to our weakest assumptions about asset pricing, they appear to be too low to be consistent with the assumption that options are redundant securities. Given their betas on the underlying index (either actual or Black–Scholes implied), at-the-money call and put options should both earn weekly returns that are two to three percent higher than their sample averages. To focus on these implications, we shift our attention to the returns on zero-beta index straddles—combinations of long positions in calls and puts that have offsetting covariances with the index. We derive an expression for zero-beta straddle returns that only requires knowledge of the beta of the call option. Using Black–Scholes call option beta values, we find that zero-beta straddles offer returns that are significantly different than the risk-free rate. Market-neutral straddles on the S&P 500 index receive average weekly returns of minus three percent per week. These findings are shown to be robust to nonsynchronous trading, mismeasurement of the call option beta, and changes in the sample period and frequency. To be certain that our results are not driven by our use of Black–Scholes betas, asymptotic normality assumptions, and discrete time returns, we test the Black–Scholes/CAPM model with the generalized method of moments (GMM). Our GMM tests confirm that straddle returns data are inconsistent with the model. Finally, to ensure that our results are not sensitive to the inclusion of the October 1987 crash, we analyze daily returns on zero-beta S&P 100 straddles from January 1986 to October 1995. Even with the crash, average returns on at-the-money S&P 100 straddles are -0.5 percent per day.

One possibility is that, in spite of including the 1987 crash, the market viewed the likelihood of additional, larger crashes during the 10-year period to be much higher than the sample frequency indicates. To confirm that the negative straddle returns are not subject to this sort of a “peso-problem,” we focus on the returns to the “crash-neutral” straddle—a straddle position with an offsetting short position in a deeply out-of-the-money put.² We find that this strategy, which has a net long position in options, but earns the same return whether the market drops 15 percent or 50 percent, still generates average losses of nearly 3 percent per week. These results indicate that op-

² This position is termed by practitioners a “ratio-call spread.”

tions are earning low returns for reasons that extend beyond their ability to provide insurance against crashes. Moreover, the results are not eliminated when transaction costs are considered. To demonstrate this, we consider the returns to the simple strategy of selling each month equal numbers of at-the-money calls and puts at the bid prices, purchasing the same number of out-of-the-money puts at the ask, and investing the remaining premium and principal in the index. Such a strategy offers a Sharpe ratio that, between 1986 and 1996, was up to twice as high as that obtained by investing in the index alone.

Because we use zero-beta straddles to establish that index option returns appear to be too low for existing models, a natural interpretation of our results is that stochastic volatility is priced in the returns of index options. In a related paper, Buraschi and Jackwerth (1999) find evidence that volatility risk is priced in options markets. Our final tests examine some evidence that stochastic volatility is priced in markets quite different from the index option market. We look at the average returns of zero-beta straddles formed with futures options on four other assets and find that straddles on assets with volatilities that are positively correlated with market volatility tend to earn negative returns. We estimate time-series regressions of CRSP's size-decile portfolio returns on the market return and a straddle return factor and find a distinct pattern in the sensitivities of the size portfolios to the straddle factor. These results offer some preliminary confirmation that volatility risk is priced in more than just options markets.

The paper proceeds as follows. In Section I, we develop a number of theoretical propositions regarding expected option returns in a setting where only systematic risks are priced. In Section II, we discuss our dataset and methods for measuring option returns. In Section III, we test the implications developed in Section I. Section IV concludes the paper.

I. Theoretical Expected Returns

Like all other assets, options must earn returns commensurate with their risk. We derive two very general results about options returns. We also discuss the implications of Black–Scholes assumptions for options returns.

A. Weak Assumptions

According to the Black–Scholes/CAPM asset-pricing framework, call options always have betas that are larger in absolute value than the assets upon which they are written. Hence, we expect call options on market indices and most individual securities to have positive long-run returns greater than those of their underlying securities. On the other hand, put options, whose returns generally hedge systematic risks, should have expected returns below the risk-free rate. We derive two theorems that produce results analogous to these in a much more general setting. To derive our results, we assume the existence of a stochastic discount factor that prices all assets according to the relation

$$E[R_i \cdot m] = 1, \quad (1)$$

where E denotes the expectation operator, R_i is the gross return of any asset, and m is the strictly positive stochastic discount factor. It is well known that such a discount factor exists in the absence of arbitrage. The stochastic discount factor is high in bad states of the world and low in good states. We make assumptions about how the stochastic discount factor is related to the returns of an option's underlying asset to derive our theorems.

As long as the stochastic discount factor is negatively correlated with the price of a given security, all call options written on that security will have expected returns above those on the security and increasing in their strike price. This is described in the following theorem.

PROPOSITION 1: *If the stochastic discount factor is negatively correlated with the price of a given security over all ranges of the security price, any call option on that security will have a positive expected return that is increasing in the strike price.*

Proof: Let the expected gross return from now until maturity on a call option with a strike price of X on an underlying security whose price has a distribution $f(y)$ be expressed as

$$E[R_c(X)] = \frac{\int_{s=X} (s - X)f(s)\partial s}{\int_{q=0} \int_{s=X} q(s - X)f(s, q)\partial s \partial q}, \quad (2)$$

where $f(y, z)$ is the joint distribution of the security price and the stochastic discount factor. The expected net return, $E[r_c(X)] = E[R_c(X)] - 1$, can be written as

$$E[r_c(X)] = \frac{\int_{s=X} (s - X)[1 - E[m|s]]f(s)\partial s}{\int_{s=X} (s - X)E[m|s]f(s)\partial s}, \quad (3)$$

where m is the stochastic discount factor. The derivative of expected net returns with respect to strike price can be expressed as

$$\frac{\int_{s=X} (s - X)f(s)\partial s \cdot \int_{s=X} E[m|s]f(s)\partial s - \int_{s=X} (s - X)E[m|s]f(s)\partial s \cdot \int_{s=X} f(s)\partial s}{\left[\int_{s=X} (s - X)E[m|s]f(s)\partial s \right]^2}. \quad (4)$$

Defining $F(s)$ as the cumulative density function that corresponds to $f(s)$ and rearranging gives

$$\frac{\partial E[r_c(X)]}{\partial X} = \frac{\int_{s=X} (s-X) \frac{f(s)}{1-F(X)} \partial s \cdot \int_{s=X} E[m|s] \frac{f(s)}{1-F(X)} \partial s - \int_{s=X} (s-X) E[m|s] \frac{f(s)}{1-F(X)} \partial s}{\left[\int_{s=X} (s-X) E[m|s] \frac{f(s)}{1-F(X)} \partial s \right]^2}. \quad (5)$$

Note that the numerator in equation (5) is simply the negative of the covariance of $s - X$ and m , conditional on the option being in the money ($s > X$):

$$\begin{aligned} -\text{Cov}[E(m|s), s - X | s > X] &= E[m | s > X] \cdot E[s - X | s > X] \\ &\quad - E[E(m|s)(s - X) | s > X]. \end{aligned} \quad (6)$$

For any m that is negatively correlated with the security price for all ranges,³ the above expression will be positive. Because a call option with a zero strike price has the same positive expected net return as that of the underlying security, the fact that net expected returns are increasing in the option strike price ($\{\partial E[r_c(X)]/\partial X > 0\}$) implies that all call options will have positive expected returns that exceed those of the underlying asset. Q.E.D.

Thus, we should expect any call options on securities whose prices are negatively correlated with the stochastic discount factor to have positive expected returns. Moreover, these returns should be increasing in the option strike price. No existing asset-pricing theory permits a stochastic discount factor that is positively correlated with the market level and most individual security prices.

The corresponding proposition for put options is as follows.

PROPOSITION 2: *If a stochastic discount factor is negatively correlated with the price of a given security over all ranges of the security price, any put option on that security will have an expected return below the risk-free rate that is increasing in the strike price.*

Proof: Let the net expected return on a put option with a strike price of X on an underlying security whose price has a distribution $f(y)$ be expressed as:

$$r_p(X) = \frac{\int_{s=X}^{s=X} (X-s)[1 - E[m|s]]f(s)\partial s}{\int_{s=X}^{s=X} (X-s)E[m|s]f(s)\partial s}, \quad (7)$$

³ Actually, the correlation must only be negative for any range of security prices for which the option is in the money ($s > X$).

where $f(x, y)$ is the joint distribution of the security price and the stochastic discount factor and m is the stochastic discount factor. After rearranging, the derivative of net returns with respect to strike price can be expressed as

$$\frac{\int_{s=X}^{s=X} (X-s)E[m|s]f(s)\partial s \cdot \int_{s=X}^{s=X} f(s)\partial s - \int_{s=X}^{s=X} (X-s)f(s)\partial s \cdot \int_{s=X}^{s=X} E[m|s]f(s)\partial s}{\left[\int_{s=X}^{s=X} (X-s)E[m|s]f(s)\partial s \right]^2}. \quad (8)$$

By the same argument as in Proposition 1, the numerator in the above equation is proportional to the covariance between $X - s$ and m , conditional on the option being in the money ($s < X$):

$$\begin{aligned} \text{Cov}[E(m|s), X-s|s < X] &= E[E(m|s)(s-X)|s < X] - E[m|s < X] \\ &\quad \cdot E[s-X|s < X]. \end{aligned} \quad (9)$$

For any m that is negatively correlated with the security price over all ranges, the above expression will be positive. Because a put option with an infinite strike price has an expected net return equal to the risk-free rate, the fact that net returns are increasing in the option strike price ($\{\partial E[r_p(X)]/\partial X > 0\}$) implies that all put options will have expected returns below the risk-free rate. Q.E.D.

Although call returns should always be positive, expected put returns can, in theory, be either positive or negative depending on their strike price. Put options that are exceedingly in the money, and have a smaller and smaller fraction of their value at expiration determined by the level of the underlying asset will have expected returns that approach the risk-free rate. Propositions 1 and 2 essentially confirm that the leverage effect in option returns should be priced regardless of whether the Black-Scholes model fits the data.

B. Black-Scholes Assumptions

Although Propositions 1 and 2 only assume the existence of a well-behaved stochastic discount factor, it is common in asset pricing to make much stronger assumptions. In particular, it is common to assume that asset prices follow Geometric Brownian Motion (GBM). It is well known that if assets follow GBM, then options are redundant securities. With a few additional assumptions, both the option pricing model of Black and Scholes (1973) and the continuous-time CAPM model of Merton (1971) hold when assets follow GBM. Expected asset returns are determined by the CAPM equation

$$E[r_i] = r + E[r_m - r]\beta_i, \quad (10)$$

where r is the short-term interest rate and r_m is the return on the market portfolio. Black and Scholes (1973) use this equation in an alternative derivation of their option pricing formula. The beta of a call option can be solved for as

$$\beta_c = \frac{s}{C} \mathcal{N} \left[\frac{\ln(s/X) + (r - \lambda + \sigma^2/2)t}{\sigma\sqrt{t}} \right] \beta_s, \quad (11)$$

where $\mathcal{N}[\cdot]$ is the cumulative normal distribution, and β_s is the beta and λ is the dividend yield of the underlying asset. If we construct an options position with a beta of zero, then its expected return should be the risk-free rate. We estimate the expected returns of zero-beta straddles below.

An alternative way to consider the risk-return implications of the Black–Scholes model is to examine the pricing kernel implied by the model. Dybvig (1981) shows that over any time interval in a Black–Scholes world, the gross (buy-and-hold) return of any asset (R_i), the gross market return (R_m), and the continuously compounded risk-free return (r) must satisfy the pricing kernel relation

$$E \left[R_i \left(\frac{R_m^{-\gamma} e^{-r}}{E[R_m^{-\gamma}]} \right) \right] = 1, \quad (12)$$

where γ is interpreted as the relative risk aversion parameter for a hypothetical consumer that holds the market portfolio. If the Black–Scholes assumptions are valid, all discrete time asset returns must satisfy this relation. We estimate γ with GMM below.

To summarize, the results above offer several testable asset-pricing implications for option returns. First, the difference between call returns and the returns on the underlying security should be positive and this difference should be greater the higher the option's strike price. Second, the difference between put option returns and the risk-free rate should be negative and this difference should be greater the lower the option's strike price. The expected returns on zero-beta straddle positions should be equal to the risk-free rate according to the Black–Scholes/CAPM paradigm. Finally, the pricing kernel that corresponds to the Black–Scholes model should price all assets.

II. Data and Method

To test the above implications, we use the Berkeley Options Database of Chicago Board Options Exchange bid-ask quotes. We examine the weekly returns of European-style call and put options on the S&P 500 index over a six-year period, from January 1990 to October 1995.⁴ We focus on the weekly

⁴ Prior to 1990, volume on the S&P 500 option was too low for the data to be useful in our tests. The database does not include out-of-the-money call options for November and December of 1995, prompting us to omit these two months from the entire study of S&P 500 options.

returns of S&P 500 options as our base case. Because our observations are weekly and are over a time period of high liquidity, we can obtain a weekly observation for each option specification studied. We also examine daily returns of American-style call and put options on the S&P 100 index over a 10-year period, from January 1986 to December 1995. Daily S&P 100 option returns are used to provide a somewhat independent verification of our results in a different market, across a longer time horizon, and at a different frequency. Moreover, in using daily observations from 1986 onward, the S&P 100 options allow us to explicitly measure how much our results are impacted if we include the crash of October 1987.

Our method for calculating option returns is as follows. For each option, we identify the first bid-ask quote after 9 a.m. Central Standard Time (CST). We take options that are to expire during the following calendar month,⁵ and therefore are roughly between 20 and 50 days to expiration. Finally, we take the midpoint of the bid-ask spread⁶ and use this to calculate weekly and daily holding-period returns for each of our options. Our weekly results only use prices observed on Tuesdays.⁷ Early in the S&P 100 sample, not all options exist on all trading days. In unreported robustness tests, we assign returns of -100 percent to all options for which we find a price on one day but no price on the next day. This leads to no discernable change in our results.

Although Sheikh and Ronn (1994), in a previous study of option returns, use logarithmic option returns, all of our calculations rely on raw net returns. Although the log-transformation of stock returns is widely accepted in the empirical asset-pricing literature, log-scaling option returns may be quite problematic. Because options held to maturity often have net returns of -1 (i.e., expire worthless), the log-transformation of any set of option returns over any finite holding period will be significantly lower than the raw net returns. Put differently, the expected log return of any option held to expiration is negative infinity. In any case, Propositions 1 and 2 are stated in terms of standard returns—not log returns.

Options are classified into five groups according to their strike prices relative to the level of the S&P at the previous close and using the same five-point increments of the exchange. We group options according to strike price as opposed to moneyness level to ensure that exactly one option exists within each group at each point in time. This, of course, comes at the cost of having the moneyness of each bin change over time. The S&P 500 and 100 index levels were taken from the underlying price accompanying the first quote on any option after 9 a.m. CST.

Finally, to examine the robustness of our results across other option markets, we examine returns of options on Treasury bond, Eurodollar, Nikkei 225 Index, and Deutsche Mark futures contracts. We calculate daily returns

⁵ These were available for all but three of the options.

⁶ We checked the spreads for outliers and removed any that significantly violated arbitrage pricing restrictions.

⁷ We refer to Wednesday whenever July 4th, December 25th, or January 1st fall on a Tuesday.

using closing prices of front-month contracts obtained from the Bridge/CRB Historical Data base from October 1988 (September 1994 for Nikkei contracts) to August 1999.

III. The Evidence

Our initial sample focuses on S&P 500 index (SPX) options over a six-year period from January 1990 through October 1995. Before turning to our results, it may be useful to review the behavior of the underlying index over this 305-week interval. Average weekly returns were 0.18067 percent or 9.84 percent on an annualized basis. The median weekly return was slightly higher at 0.1998 percent, with a minimum of -5.39 percent and maximum of 6.05 percent. Overall, the returns of the S&P 500 during this period can be viewed as fairly representative of the underlying distribution. They do not contain a crash of magnitude similar to that of October 1987, nor do they focus on a period of abnormally high returns (i.e., 1995 to 1999). The S&P 100 (OEX) time series includes the 10-year period of January 1986 to December 1995. Average returns were 0.04 percent per day, or 12.68 percent on an annualized basis. Median daily returns were 0.05 percent, with a minimum return of -21.21 percent and a maximum of 8.91 percent. Clearly, a major difference between the two series is that the OEX data include the 1987 stock market crash and the SPX data do not.

A. Average Option Returns

Tables I and II record a variety of statistics for weekly SPX and daily OEX option returns over 6- and 10-year sample periods. We record the mean, median, minimum, and maximum weekly returns for each of the five groups of strike prices. The groups range from options with strike prices 10 to 15 points below the index level to those with strike prices 5 to 10 points above the index. The t -statistic associated with a null hypothesis of zero mean return is recorded in the second row. We also report the average Black-Scholes implied beta for each of the strike price ranges.

Looking at mean returns of S&P 500 call options in Table I, we see that they do indeed earn positive average returns in excess of those on the underlying security. An at-the-money call option tends to earn an average return of between 1.85 percent and 2.00 percent per week, or between 96 percent and 110 percent on an annualized basis. Moreover, average returns are increasing (almost) monotonically in the strike price. Call options that are between 5 and 10 points out of the money earn 2.94 percent more per week (153 percent annualized) than those equivalently in the money. Although these differences in returns across strike prices are not statistically significant, they are qualitatively consistent with Proposition 1. Not surprisingly, the median, minimum, and maximum statistics demonstrate a substantial degree of positive skewness in the call returns, which is increasing in their strike prices. S&P 100 call option returns exhibit similar characteristics.

Table I
Call Option Returns

This table reports summary statistics for call option returns. The SPX sample period is January 1990 to October 1995 (305 weeks). The OEX sample period is January 1986 to October 1995 (2519 days). Bid-ask spreads are checked for significant outliers. SPX returns are recorded in weekly percentage terms. OEX returns are recorded in daily percentage terms. $X - s$ denotes the difference between the option's strike price and the underlying price. BS β denotes the option beta calculated from the Black-Scholes formula.

$X - s$	-15 to -10	-10 to -5	-5 to 0	0 to 5	5 to 10
Panel A: Weekly SPX Call Option Returns					
Mean return	1.48	1.19	1.85	2.00	4.13
<i>t</i> -Statistic	(0.79)	(0.53)	(0.66)	(0.55)	(0.85)
Median	0.0	-1.99	-4.46	-9.55	-17.39
Minimum	-80.67	-86.51	-89.33	-92.85	-92.31
Maximum	141.82	190.24	256.63	426.65	619.41
Mean BS β	21.14	25.23	31.20	40.02	55.72
Panel B: Daily OEX Call Option Returns					
Mean return	0.62	0.55	0.67	0.80	0.85
<i>t</i> -Statistic	(1.97)	(1.61)	(1.61)	(1.55)	(1.33)
Median	0.18	0	0	-1.83	-3.28
Minimum	-69.24	-75.16	-73.39	-75.36	-75.91
Maximum	321.95	91.89	118.47	190.35	267.15
Mean BS β	17.54	20.95	25.91	33.62	46.52

Here, returns are also in excess of those of the underlying security—between 0.62 percent and 0.85 percent per day—are somewhat statistically significant, and are increasing in the strike price.

Table I also includes the average Black-Scholes beta for each set of call options. We calculate call betas with the Black-Scholes formula for call betas given in equation (11). For option volatility, we use the opening level of the CBOE's S&P 100 option implied volatility index (VIX). For the dividend yield, we use the average annual yield calculated from daily realized dividends on the S&P 100 index during the life of each option. The dividend yield and volatility of the S&P 100 index should be nearly identical to those of the S&P 500 index. The Black-Scholes betas of the call options are, as expected, increasing across strike prices. The magnitudes are, indeed, quite large, ranging from 21.14 to 55.72 for the SPX and 17.54 to 46.52 for the OEX. Unreported time-series regressions of option returns on underlying index returns generate beta estimates that are quite similar in magnitude to the Black-Scholes betas. Although the call options appear to earn impressive returns, a quick glance at their corresponding betas suggests that the returns are actually far too low to be consistent with the Black-Scholes/CAPM model. Assuming a market risk premium of six percent, the at-the-money option returns should be around four percent per week, making both

Table II
Put Option Returns

This table reports summary statistics for put option returns. The SPX sample period is January 1990 to October 1995 (305 weeks). The OEX sample period is January 1986 to October 1995 (2519 days). Bid-ask spreads are checked for significant outliers. SPX returns are recorded in weekly percentage terms. OEX returns are recorded in daily percentage terms. $X - s$ denotes the difference between the option's strike price and the underlying price. BS β denotes the option beta calculated from the Black-Scholes formula.

$X - s$	-15 to -10	-10 to -5	-5 to 0	0 to 5	5 to 10
Panel A: Weekly SPX Put Option Returns					
Mean return	-14.56	-12.78	-9.50	-7.71	-6.16
<i>t</i> -Statistic	(-4.22)	(-3.95)	(-3.05)	(-2.81)	(-2.54)
Median	-28.27	-24.96	-21.57	-15.06	-11.34
Minimum	-84.03	-84.72	-87.72	-88.90	-85.98
Maximum	475.88	359.18	307.88	228.57	174.70
Mean BS β	-36.85	-37.53	-35.23	-31.11	-26.53
Panel B: Daily OEX Put Option Returns					
Mean return	-2.30	-2.02	-1.79	-1.42	-1.22
<i>t</i> -Statistic	(-2.92)	(-2.88)	(-3.06)	(-2.88)	(-2.98)
Median	-7.56	-6.55	-5.48	-4.10	-3.02
Minimum	-73.11	-68.52	-72.74	-67.31	-67.39
Maximum	795.52	759.73	527.53	394.15	300.00
Mean BS β	-32.92	-31.87	-29.17	-25.33	-21.31

SPX and OEX call returns roughly two percent too low. Given their substantial comovement with the underlying index, call options might be expected to earn far greater returns than they do.

Turning to put option returns in Table II, we see results that are consistent with Proposition 2. Put options have returns that are both economically and statistically negative. At-the-money put options lose between 9.5 percent and 7.7 percent per week. As claimed by Proposition 2, puts have expected returns that are increasing in their strike price. The mean returns increase monotonically in the strike price, ranging from -14.56 percent for the most out-of-the-money options to -6.16 percent for those deeply in the money. Again, as expected, the median, minimum, and maximum statistics indicate that put returns exhibit substantial positive skewness. Daily OEX put returns are also consistent with Proposition 2. Put options lose between 1.22 percent and 2.30 percent per day, and higher strike prices exhibit uniformly smaller average losses.

The Black-Scholes betas for put options are large and negative, ranging from -36.85 to -26.53 for the SPX and from -32.92 to -21.31 for the OEX, and increasing across strike prices. Again, however, mean returns appear to be too low to be consistent with the betas. With a market risk premium of six percent per year, at-the-money put options should earn weekly returns of

around minus four percent according to their betas. With average at-the-money SPX and OEX put returns averaging around minus eight percent, put returns are also too low to be consistent with the Black–Scholes/CAPM model. Hence, it seems that although option returns are qualitatively consistent with Propositions 1 and 2 and consistent with the appropriate pricing of the leverage effect, back-of-the-envelope calculations indicate that both calls and puts earn returns that are too low. In the following sections, we investigate these issues in greater detail.

B. Straddle Returns

To determine whether call and put options earn returns that are below those justified by the CAPM and the assumption that they are redundant securities, we direct our attention toward the returns of straddle positions. By forming zero-beta, or zero-delta, straddle positions by combining puts and calls with the same strike price and maturity in proportion to their betas, we can abstract from the leverage effect and focus on the pricing of higher moments of security returns. According to the Black–Scholes/CAPM model, zero-beta straddles should have expected returns equal to the risk-free rate.

Although straddle returns are not sensitive to market returns, they are sensitive to market volatility. When volatility is higher than expected, straddles have positive returns. When volatility turns out lower than expected, straddles have negative returns. In a factor model framework, straddles have large, positive volatility betas. They allow investors to hedge volatility risk. This makes straddle returns ideal for studying the effects of stochastic volatility. Because investors presumably dislike volatile states of the world, assets with negative volatility betas or assets that lose value when volatility increases are relatively risky assets. Likewise, assets with positive volatility betas, like straddles, are relatively low-risk assets. In fact, because straddles insure investors against volatile states of the world, if there is a risk premium for volatility risk, expected zero-beta straddle returns will be less than the risk-free rate.

Specifically, we form zero-beta straddles by solving the problem

$$\begin{aligned} r_v &= \theta r_c + (1 - \theta) r_p \\ \theta \beta_c + (1 - \theta) \beta_p &= 0, \end{aligned} \tag{13}$$

where r_v is the straddle return, θ is the fraction of the straddle's value in call options, and β_c and β_p are the market betas of the call and put, respectively. This problem is solved by the weight function

$$\theta = \frac{-\beta_p}{\beta_c - \beta_p}. \tag{14}$$

As this equation demonstrates, forming zero-beta straddles requires knowing the betas of call and put options. Fortunately, as long as put-call parity holds, knowing the beta of a call option is sufficient to determine the beta of the put option with the same strike price and maturity. Put-call parity states that, at any time t ,

$$C(t) = P(t) + s(t) - (X + D)e^{-r(T-t)}, \quad (15)$$

where C is the price of a call, P is the price of a put, s is the price of the underlying asset, D is the amount of dividends to be paid by the stock at the maturity of the options, and $r(T - t)$ is the risk-free rate multiplied by the time to maturity. Dividing the put-call parity condition at time $(t + \Delta t)$, where Δt is any quantity less than $T - t$, by $C(t)$ gives

$$\frac{C(t + \Delta t)}{C(t)} = \frac{P(t)}{C(t)} \frac{P(t + \Delta t)}{P(t)} + \frac{s(t)}{C(t)} \frac{s(t + \Delta t)}{s(t)} + \frac{(X + D)e^{-rt}}{C(t)}, \quad (16)$$

which implies that

$$\beta_c = \frac{P}{C} \beta_p + \frac{s}{C} \beta_s. \quad (17)$$

Assuming that $\beta_s = 1$, we can combine equations (14) and (17) to define the return to a zero-beta index straddle as

$$r_v = \frac{-C\beta_c + s}{P\beta_c - C\beta_c + s} r_c + \frac{P\beta_c}{P\beta_c - C\beta_c + s} r_p. \quad (18)$$

The only term in equation (17) that is not directly observable is the call option's beta, β_c . Thus, the ambiguity about the straddle weights that achieve a beta of zero can be reduced to just one parameter. In the empirical work below, we use the Black-Scholes beta to identify this parameter. We define the Black-Scholes beta as in equation (11).

In Table III we measure the weekly returns of S&P 500 and daily returns of S&P 100 zero-beta straddle positions across strike prices. The S&P 500 straddles are rebalanced each week and the S&P 100 straddles are rebalanced each day. As before, we record the mean, median, minimum, and maximum returns on the positions. Average returns of zero-beta straddle positions are invariably negative. At-the-money SPX straddles have average returns of -3.15 percent per week and all the SPX straddles we examine lose between 2.89 percent and 4.49 percent per week, with lower strike price straddles losing the most. All of the negative returns are highly statistically significant. The precision of our zero-beta straddle calculations depends critically on our method used to estimate call option betas. However, we allow for some measurement error in our beta calculations by averaging straddle

Table III
Returns of Zero-Beta Straddles

The table reports summary statistics for the returns of zero-beta straddles. The SPX sample period is January 1990 to October 1995 (305 weeks). The OEX sample period is January 1986 to October 1995 (2519 days). Bid-ask spreads are checked for significant outliers. SPX returns are recorded in weekly percentage terms. OEX returns are recorded in daily percentage terms. The straddle return is calculated according to equation (18) where β_c is the Black-Scholes beta of the call option, calculated using the CBOE's VIX index as implied volatility. $X - s$ denotes the difference between the option's strike price and the underlying price. The t -statistics test the null hypothesis that mean zero-beta straddle returns are zero.

$X - s$	-15 to -10	-10 to -5	-5 to 0	0 to 5	5 to 10
Panel A: Weekly SPX Straddle Returns					
Mean return	-4.49	-4.28	-3.15	-3.15	-2.89
t -Statistic	(-5.44)	(-4.75)	(-2.89)	(-2.72)	(-2.38)
Median	-7.17	-7.96	-7.21	-8.27	-8.17
Minimum	-29.57	-38.24	-41.46	-35.39	-44.43
Maximum	124.29	102.53	115.06	131.71	145.84
Setting $\beta_c = 0.9$ (Black-Scholes β_c)					
Mean return	-2.96	-3.23	-2.42	-2.63	-2.49
t -Statistic	(-3.48)	(-3.35)	(-2.14)	(-2.16)	(-1.98)
Setting $\beta_c = 1.1$ (Black-Scholes β_c)					
Mean return	-4.30	-4.17	-3.07	-3.10	-2.85
t -Statistic	(-5.29)	(-4.62)	(-2.82)	(-2.66)	(-2.35)
Panel B: Daily OEX Straddle Returns					
Mean return	-0.57	-0.50	-0.50	-0.48	-0.65
t -Statistic	(-3.01)	(-2.51)	(-2.76)	(-2.84)	(-3.84)
Median	-1.47	-1.37	-1.47	-1.51	-1.78
Minimum	-35.02	-31.35	-39.04	-40.44	-43.20
Maximum	253.19	279.58	233.55	206.41	152.76
Setting $\beta_c = 0.9$ (Black-Scholes β_c)					
Mean return	-0.35	-0.41	-0.44	-0.44	-0.67
t -Statistic	(-2.18)	(-2.40)	(-2.74)	(-2.76)	(-4.26)
Setting $\beta_c = 1.1$ (Black-Scholes β_c)					
Mean return	-1.38	-0.93	-0.81	-0.68	-0.83
t -Statistic	(-4.95)	(-4.16)	(-4.20)	(-3.74)	(-4.70)
Setting In-Sample Beta to Zero					
Mean return	-0.39	-0.33	-0.35	-0.32	-0.51
t -Statistic	(-2.31)	(-1.86)	(-2.12)	(-2.01)	(-3.18)
Dropping Puts and Calls more than 15 min Apart					
Mean return	-0.46	-0.44	-0.49	-0.44	-0.58
t -Statistic	(-2.05)	(-2.11)	(-2.72)	(-2.59)	(-3.44)
Subsample: 1986-1990					
Mean return	-0.33	-0.14	-0.27	-0.26	-0.56
t -Statistic	(-0.95)	(-0.36)	(-0.82)	(-0.89)	(-1.99)
Subsample: 1991-1995					
Mean return	-0.81	-0.86	-0.72	-0.69	-0.74
t -Statistic	(-4.89)	(-6.15)	(-4.81)	(-4.16)	(-3.86)

returns for both the original Black–Scholes betas and call betas that are 10 percent above and below the Black–Scholes betas. As we can see, even after allowing for this measurement error, most of the returns remain significantly negative.

The SPX zero-beta straddle results are subject to criticism, however, on the grounds that the time period considered does not include a crash of the magnitude of the crash of October, 1987. Because straddle positions deliver substantial payoffs when the market crashes, it is important to recognize the probability and magnitude of these payoffs in measuring their expected returns. Indeed, the fact that SPX straddles with out-of-the-money put options lose more than those with out-of-the-money calls suggests that the negative straddle returns may be related to their ability to hedge crashes. To address this, we also calculate straddle returns on the S&P 100 index from January 1986 to October 1995. Because we calculate daily rather than weekly returns, we ensure that the full effect of the crash is captured.

In the second half of Table III, we see that daily OEX straddle returns are substantially negative even after including the crash of 1987. OEX straddles lose between 0.48 percent and 0.65 percent per day, which is comparable to the weekly returns of SPX straddles. Dividing the sample into two subperiods does not change our results, though some strike price ranges lose significance. Like the S&P 500 straddle returns, OEX straddle returns are robust to measurement errors in the betas of calls. As an additional robustness check, we set the theoretical position beta in equation (13) to a constant such that the in-sample straddle beta is exactly equal to zero. This produces straddle returns that remain significantly negative and are roughly equivalent to those obtained from a 10 percent reduction in the Black–Scholes call betas. Unlike S&P 500 options, OEX options are not always very liquid. To be sure that nonsynchronous trading does not drive our results, we drop all put/call observations that are not within 15 minutes of each other, but this has no discernible effect on our results.

These results strongly suggest that some risk other than market beta risk is being priced in option returns. Investment strategies that exploit the apparent overpricing of calls and puts need not assume high degrees of market risk. By simultaneously shorting both options—by selling straddles—investors can earn excess returns of around three percent per week, without exposing themselves to significant market risk.

C. GMM Tests

Although the zero-beta straddle results reported in the previous section are compelling, they are subject to three important criticisms. First, the results can be considered tests of the continuous-time CAPM, which is expressed in terms of instantaneous returns. We use daily and weekly buy-and-hold returns rather than instantaneous returns. Second, the results use the Black–Scholes model to calculate betas, or risk exposures, even though they reject the Black–Scholes model. Moreover, the Black–Scholes betas that

we set to zero are betas for instantaneous returns. Returns measured over a longer horizon may have betas that systematically differ from Black–Scholes betas. Third, although it is plausible that straddle returns are independent of each other through time, they are probably not identically distributed. It is not clear that the means we report are asymptotically normally distributed with the simple standard deviations that we calculate.

We address these criticisms by employing a more sophisticated test of the Black–Scholes model. As stated above, Dybvig (1981) shows that the Black–Scholes model implies a pricing kernel of the form in equation (12),

$$E \left[R_i \left(\frac{R_m^{-\gamma} e^{-r}}{E[R_m^{-\gamma}]} \right) \right] = 1.$$

This pricing kernel is valid for discrete time returns measured at any frequency. Furthermore, because it does not explicitly rely on any beta calculations, equation (12) can be estimated with any asset returns series, regardless of the asset betas. It must hold even if instantaneous betas change systematically through time. To obtain as much statistical power as possible, we estimate equation (12) with the zero-beta straddle returns described in the previous section. However, the moment condition must hold even if there is mismeasurement in the betas that we calculate to form our straddles.

We estimate γ in equation (12) with GMM. We estimate 10 different parameter values, using market returns and zero-beta straddle returns of a particular moneyness as the returns series for each estimation. Because each estimation calculates one parameter estimate with two returns series, each γ estimate has an accompanying test of overidentifying restrictions. This test of overidentifying restrictions is a convenient way to test the Black–Scholes/CAPM model. Because GMM makes very weak assumptions about the statistical properties of the model (i.e., it does not require observations to be identically distributed), we can be confident that the tests are well specified. We report parameter estimates and test statistics in Table IV.

Panel A of Table IV reports GMM estimates of γ and test statistics using market returns and the weekly SPX straddle returns described in Table III. The estimates of γ range from 5.68 to -6.68 , and they appear to be roughly declining in $(X - s)$. Because γ can be interpreted as a risk aversion coefficient, negative estimates of γ do not make sense. The (unreported) standard errors of the γ estimates are quite large. More significantly, the test of overidentifying restrictions rejects each of the models proposed. This is strong evidence that the Black–Scholes/CAPM model does not describe straddle and market returns well.

To be certain that negative average straddle returns are driving the rejections of equation (12), we estimate the model with straddle returns adjusted so that their sample averages are set to zero. From all straddle returns, we subtract their sample average so that the average adjusted return is zero. For example, we estimate the model for the moneyness category $(-15$

Table IV
GMM Tests of the Black-Scholes/CAPM Model

This table reports the results of GMM tests of the Black-Scholes moment condition for the discrete time returns of any asset,

$$E \left[R_i \left(\frac{R_m^{-\gamma} e^{-r}}{E[R_m^{-\gamma}]} \right) \right] - 1 = 0.$$

The estimates reported in the table are calculated with weekly returns of the S&P 500 index and of zero-beta, SPX straddle positions. The chi-square test statistics reported are for the test of overidentifying restrictions. $X - s$ denotes the difference between the option's strike price and the underlying price. The sample period is January 1990 to October 1995 (305 weeks). Bid-ask spreads are checked for significant outliers. Panel A gives estimation results for the original data, described further in Table III. Panel B gives estimation results when the sample mean of all straddle returns is set to zero.

$X - s$	-15 to -10	-10 to -5	-5 to 0	0 to 5	5 to 10
Panel A: Original Data					
$\hat{\gamma}$	5.41	0.55	5.68	-0.09	-6.68
Chi-square	12.36	12.38	7.90	5.39	3.91
p-Value	0.0004	0.0004	0.0049	0.0202	0.0478
Panel B: Setting Average Straddle Returns to Zero					
$\hat{\gamma}$	6.09	5.58	4.76	5.09	4.00
Chi-square	0.09	0.45	0.13	0.04	0.10
p-Value	0.7697	0.8314	0.7145	0.8350	0.7543

to -10) with market returns and with 4.49 percent added to all straddle returns. The results are reported in Panel B. Adjusting the means of the straddle returns appears to allow the model to fit the data. Whereas equation (12) is strongly rejected for each original straddle returns series, it is not rejected for any of the zero-mean straddle returns series. Furthermore, the estimates of γ produced by forcing average straddle returns to zero range from 4.00 to 6.09, a very reasonable range for risk-aversion coefficients. This indicates that the negative straddle returns reported in Table III are, in fact, inconsistent with the Black-Scholes/CAPM model.

D. Average Returns of "Crash-Neutral Straddles"

The daily returns on OEX straddles indicate that including a single crash of the magnitude of that of October 1987 does not eliminate the substantial negative expected returns to straddle positions. However, it is still possible that, during the 10-year period analyzed, the market viewed the likelihood of additional, larger crashes to be high. Indeed, Jackwerth and Rubinstein (1996) argue that the market-assigned probability of a three to four standard deviation crash increased 10-fold following the 1987 crash. Because

these expectations did not materialize, straddle positions may have earned returns that were significantly more negative than expected. An alternative possibility is that the market may simply have priced crash risk exceedingly high during this period. Securities that delivered payoffs in crash states had their expected returns set at very low levels by the market.

To determine whether one of these conjectures accounts for the observed low returns to straddle positions, we analyze the returns of “crash-neutral straddles,” or what practitioners term “ratio-call spreads.” The zero-beta straddle positions that we analyze in Table III have very large payoffs, and hence very large returns when the market crashes. However, a crash-neutral straddle’s return during a market crash is limited to some level that is specified when the position is created. For example, a crash-neutral straddle can be constructed so that its payoff increases as the market loses up to 10 percent of its value, but its payoff remains the same whether the market loses 10 percent or 50 percent of its value. The position can be achieved by purchasing a straddle position and selling a deeply out-of-the-money put option.⁸ In this way, measures of the position’s expected returns are not downward-biased by infrequent crash observations or high-priced crash risk. Because the position takes in premium from the sale of the out-of-the-money put, crash probabilities and insurance costs are explicitly incorporated into the position’s return up-front.

We calculate zero-beta, crash-neutral straddles by solving for the fraction of the portfolio allocated to the call position when the remaining weight is placed in offsetting long and short positions of near- and out-of-the-money puts, respectively. To do so, we adjust equation (13) as follows:

$$\begin{aligned} r_v &= \theta r_c + (1 - \theta) r_{p^*} \\ \theta \beta_c + (1 - \theta) \beta_{p^*} &= 0, \end{aligned} \tag{19}$$

where r_{p^*} and β_{p^*} denote the return and beta of the put leg of the straddle net an offsetting short position in an out-of-the-money put option. Because the long and short put positions are identical in terms of number of option contracts, instead of dollar amounts, the total put position beta can be calculated as

$$\beta_{p^*} = \frac{\partial p^*}{\partial s} \frac{s}{p^*} = (\Delta_{p_1} - \Delta_{p_2}) \cdot \frac{s}{p_1 - p_2}, \tag{20}$$

where p_1 is the price of the put option leg of the straddle, p_2 is the price of the out-of-the-money put option, and Δ_{p_1} and Δ_{p_2} are the respective option

⁸ Alternatively, the position can be created by selling deeply in-the-money call options and purchasing a larger number of call options at a higher strike price. Because deeply out-of-the-money put options trade far more frequently than deeply in-the-money calls, we use the former construction.

deltas. Following the arguments of Section B, the zero-beta weight can be written as

$$\theta = \frac{-\beta_{p^*}}{\beta_c - \beta_{p^*}} = \frac{\Delta_{p_2} - \Delta_{p_1}}{\Delta_c \frac{p_1 - p_2}{c} + \Delta_{p_2} - \Delta_{p_1}}. \quad (21)$$

Here, we can only use put-call parity to solve for the beta of the straddle put as a function of that of the straddle call. Thus solving for a zero-beta, crash-neutral straddle requires two parameters: the call option's beta, and the out-of-the-money put option's beta. Again, we rely on Black–Scholes to provide the beta estimates, though, as before, the results are quite robust to mismeasurement.

When implementing the crash-neutral straddle, one must also select the point at which the position's returns become invariant to further declines in the index—that is, one must select the strike price for the out-of-the-money put option. In our tests, we select the put option with the lowest available strike price that is at most 15 percent below the index level. The average moneyness of the out-of-the-money puts used in our results is 89 percent. A typical zero-beta, crash-neutral straddle's payoff is depicted in Figure 1. Here, we show what the payoff to a \$100 investment in a zero-beta, crash-neutral straddle would have been on January 3, 1995, when the index was at 458.14. Using the option prices, their Black–Scholes implied betas, and equation (21), we implement the \$100 straddle position by purchasing 8.32 call contracts with a strike price of 465, purchasing 6.26 put options with a strike of 465, and selling 6.26 put options with a strike of 400. As we can see, the strategy has its payoff capped at around \$400 upon a large decline in the index level.

Table V shows that daily S&P 100 and weekly S&P 500 zero-beta straddle returns remain substantially negative even if we sell off their crash insurance. Average OEX returns are -1.02 percent per day whereas SPX returns are -3.24 percent per week. As we can see, selling the crash insurance substantially lowers the maximum and overall returns on the OEX options, presumably because this capped their returns during the 1987 crash. Unreported tests verify that our results are robust to mismeasurement of the call beta or the out-of-the-money put beta.

One possible concern with our effort to explicitly capture the 1987 crash by using daily returns is that quotes for out-of-the-money put options may not have been available as the market was declining on Monday the 19th. To address this, we use transaction prices for out-of-the-money put options whenever a quote is unavailable. The performance of the out-of-the-money put option is recorded in Panel C of Table V. Each price reflects the first transaction to have occurred after 9 a.m. As we can see, on the 19th and 20th of October, the first transactions in out-of-the-money put options did not occur until 9:50 and 10:53, respectively. Nonetheless, if we neutralize the straddle's crash exposure at these prices, the returns of the overall position, which is short the out-of-the-money put option, are only made worse.




Figure 1. \$100 crash-neutral, zero-beta straddle payoff: January 3, 1995. This figure depicts the payoff diagram for the “crash-neutral” straddle position that our algorithm specifies for January 3, 1995. The horizontal axis measures option moneyness, or the ratio of the spot price to the strike price of the position. The vertical axis measures the position’s payoff at maturity.

Finally, to see how robust our findings are to the inclusion of transaction costs, and to see how well a strategy that exploits the apparent overpricing of options would fare in “real time,” we consider the following strategy. Each month, on the day following expiration, we sell (at the bid price) an equal number of at-the-money calls and puts that are now one month to maturity. At the same time, we purchase (at the ask) an equivalent number of out-of-the-money puts to hedge against large crashes. Lastly, we invest the remaining proceeds and the principal in the underlying index. Table VI reports the monthly returns to this strategy from 1986 through 1995 in the OEX markets for a range of option premium levels. Over this 10-year period, the strategy beat the index soundly. At a risk-free rate of five percent per annum, the strategy offers a Sharpe ratio of up to 31.31 percent—more than double that of the index alone (14.47 percent). Moreover, the portfolio can consist of up to 15 percent of option premium without incurring maximum monthly losses of greater than 50 percent. To see whether the strategy’s high Sharpe ratio is a statistically significant improvement over that of the market, we employ a Britten-Jones (1999) *F*-test to test the hypothesis that the mean-variance efficient portfolio contains no exposure to straddles. The *F*-statistic associated with the restriction of a zero straddle weight is 8.76 with a *p*-value of 0.0037. Indeed, even with the incorporation of transaction

Table V
Returns of Zero-Beta, Crash-Neutral Straddles

This table reports summary statistics on the returns of zero-beta, crash-neutral straddles on the OEX and SPX. The SPX sample period is January 1990 to October 1995 (305 weeks). The OEX sample period is January 1986 to October 1995 (2519 days). Both straddles are neutral to crashes of larger than 15 percent. Bid-ask spreads are checked for significant outliers. Returns are recorded in daily and weekly percentage terms. The straddle return is calculated according to equation (21) where β_c is the Black–Scholes beta of the call option, and β_{p*} is that of the combined put position. Both are calculated using the CBOE's VIX index as implied volatility. X/s denotes the option's strike price divided by the underlying price. The t -statistics test the null hypothesis that mean zero-beta straddle returns are zero.

Panel A: Weekly SPX Straddle Returns					
Mean return		−3.24			
<i>t</i> -Statistic		(−2.15)			
Median		−7.51			
Minimum		−210.41			
Maximum		149.15			
Panel B: Daily OEX Straddle Returns					
Mean return		−1.02			
<i>t</i> -Statistic		(−4.09)			
Median		−1.50			
Minimum		−236.10			
Maximum		102.79			
Panel C: Out-of-the-Money Put Data Around October 19, 1987					
Date	Time	Strike	<i>X/s</i>	Price	Return
10/16	9:04	280	0.965	3.75	780.0
10/19	9:50	255	0.931	22.00	213.6
10/20	10:53	200	0.925	80.00	−72.5
10/21	9:02	200	0.858	22.00	−36.4
10/22	9:56	220	0.867	22.00	−18.2

costs, a strategy of selling crash-neutral straddles offers returns that are highly significant—both economically and statistically.

E. Other Markets

To investigate whether our results are a special feature of the S&P index option markets or whether they apply more broadly, we examine average straddle returns across a variety futures options. Specifically, we measure the daily returns of straddle positions in the U.S. Treasury bond, Eurodollar, Nikkei 225 Index, and Deutsche Mark futures options contracts. We calculate returns of zero-beta positions⁹ by setting the weights according to equa-

⁹ Our positions have betas that are zero with respect to the underlying contract—not with respect to the market portfolio.

Table VI
Returns from Selling Straddles Versus Index

This table compares the monthly OEX index returns to those from a strategy of selling at-the-money straddles, purchasing an offsetting 15 percent out-of-the-money put, and investing the remaining premium and principal in the index. The sample period is January 1986 to October 1995 (118 months). Bid-ask spreads are checked for significant outliers. The at-the-money options are sold at the bid price. The out-of-the-money put is purchased at the ask price. All options are purchased the day following expiration, mature the following month, and are held to maturity. Returns are recorded in monthly percentage terms.

Option Premium (percent)	0	5	10	15	20
Mean return	0.99	1.94	2.89	3.84	4.79
Median	1.32	3.96	4.85	6.98	8.48
Minimum	-21.53	-27.22	-35.00	-47.45	-59.90
Maximum	11.05	7.01	11.93	16.98	22.03
Std. deviation	3.99	5.48	8.04	10.96	14.01
Sharpe ratio	14.47	27.91*	30.82*	31.29*	31.27*

* The Britten-Jones (1999) *F*-statistic is 8.76 (*p*-value of 0.004) on the restriction that the straddle weight in the mean-variance efficient portfolio is zero.

tion (14) and using the average of at-the-money put and call volatilities implied by Black-Scholes. Our results are presented in Table VII.

Consistent with the S&P index option returns, average call returns are positive across all four markets and put returns are negative for all but the Nikkei contract. Average straddle returns are negative for options on Treasury bond, Eurodollar, and Nikkei index futures. Treasury bond straddle returns, at -0.69 percent per day, are as statistically and economically significant as those of the S&P index options. Returns on Eurodollar and Nikkei straddles, although economically large and negative at -0.33 percent and -0.29 percent per day, respectively, are insignificant statistically. Deutsche Mark straddle returns are virtually zero.

Although average straddle returns in other markets are not uniformly negative like index straddle returns, they are not necessarily inconsistent with the conjecture that volatility risk is priced. The central intuition of asset-pricing theory is that only undiversifiable risks are priced in equilibrium. To examine whether volatility risk is priced, we need to be specific about which volatility is systematic. If the only systematic volatility in the economy is market volatility, then only assets with volatilities that are positively correlated with that of the market should earn a risk premium. Alternatively, if several volatility factors are priced, then we should not expect straddle returns to be consistently related to the correlation of their underlying asset's volatility with the market's volatility.

To investigate this further, in Panel B of Table VII we present correlations between the implied volatilities of our four markets and the VIX. Assuming that the VIX represents an approximate measure of the volatility of the

Table VII
Zero-Beta Straddle Returns of Other Assets

This table reports average zero-beta straddle returns for Treasury bond (TB), Eurodollar (ED), Nikkei 225 Index (NI), and German Mark (DM) futures options. It also reports correlation coefficients for the implied volatilities of these options (σ) and the VIX. The data consist of daily closing prices obtained from Bridge/CRB Historical Database for the time period October 1988 (September 1994 for Nikkei) to August 1999.

Panel A: Average Daily Returns				
Underlying	T-bonds	Euro\$	Nikkei	D-Marks
Mean call return	0.53	1.20	0.12	0.12
Mean put return	-1.51	-1.52	0.43	-0.03
Mean strad return	-0.69	-0.33	-0.29	0.01
<i>t</i> -Statistic	(-5.04)	(-1.68)	(-0.57)	(0.03)
Number of days	1783	2580	1075	2687
Panel B: Implied Volatility Correlations				
	σ_{TB}	σ_{ED}	σ_{NI}	σ_{DM}
σ_{ED}	0.25**			
σ_{NI}	-0.06	-0.07*		
σ_{DM}	0.01	0.13**	0.14**	
VIX	0.12**	-0.12**	0.08*	-0.16**

*, ** denote statistical significance at the 0.05 and 0.0001 levels, respectively.

market portfolio, we see some evidence that market volatility is priced. Our highest correlation is between the Treasuries and the VIX at 12 percent, which is consistent with the highly significant negative returns on Treasury bond straddles. Our lowest correlation is between the Deutsche Mark and the VIX at -16 percent, which is consistent with the relatively high DM straddle returns. At eight percent, the correlation between Nikkei implied volatility and the VIX is also fairly positive, accounting for the somewhat negative Nikkei returns. On the other hand, the Eurodollar volatility correlation is -12 percent, which is inconsistent with their -0.33 percent daily straddle return. All of the correlations with VIX in Panel B are statistically significant at conventional levels.

We interpret these results as very tentative evidence that market volatility risk is priced. However, it seems likely that other volatilities, such as interest rate, exchange rate, or production volatility are both undiversifiable and important to investors. Our evidence does not preclude the possibility that several types of volatility risk are independently priced.

F. Time-series Returns Regressions

The negative returns of straddle positions, particularly on securities whose volatilities are positively correlated with that of the market, indicates that a

Table VIII
Time-series Returns Regressions

This table contains the results of time-series regressions that test whether straddle returns are priced as a factor. The dependent variable of each regression is the excess return of one of CRSP's size-decile portfolios, and each regression is estimated with monthly data from January 1986 through December 1995 (120 months). The intercepts are expressed in terms of percent per month. Zero-beta straddle returns, r_{vt} , are calculated by estimating at-the-money returns with a weighted average of daily zero-beta straddle returns, and cumulating these daily returns each month. The GRS F -test reported at the bottom of the table is from Gibbons, Ross, and Shanken (1989). r_{mt} is the return of CRSP's value-weighted index of all NYSE and AMEX stocks and r_f is CRSP's short-term Treasury bill rate.

$$r_{it} - r_f = \alpha_i + \beta_{im}(r_{mt} - r_f) + \beta_{iv}(r_{vt} - r_f) + \varepsilon_{it}$$

$r_i - r_f$	α_i	t -Statistic	β_{im}	t -Statistic	β_{iv}	t -Statistic	Adjusted R^2
Size 1	-0.287	-0.53	0.84	6.31	-0.0328	-2.08	0.33
Size 2	-0.523	-1.51	0.84	9.96	-0.0296	-2.96	0.55
Size 3	-0.600	-2.32	0.83	13.19	-0.0301	-4.03	0.69
Size 4	-0.627	-2.93	0.86	16.45	-0.0291	-4.70	0.77
Size 5	-0.582	-3.08	0.89	19.29	-0.0242	-4.42	0.81
Size 6	-0.353	-2.02	0.93	21.85	-0.0237	-4.70	0.85
Size 7	-0.273	-1.75	0.99	26.10	-0.0189	-4.21	0.88
Size 8	-0.202	-1.73	0.96	33.66	-0.0156	-4.62	0.92
Size 9	-0.096	-0.97	1.04	42.98	-0.0090	-3.16	0.95
Size 10	0.094	2.19	1.00	95.98	0.0059	4.81	0.99

GRS F -Test = 2.31 ($p = 0.0168$)

market volatility risk factor may be important in explaining security returns. To pursue this, our final set of results focus on time-series regressions of portfolio returns on market return and volatility factors. Specifically, we regress excess returns of CRSP's size-decile portfolios on excess returns of the market and the excess returns of our zero-beta straddles. To calculate our straddle returns, we first estimate a daily at-the-money zero-beta straddle return by taking a weighted average of the zero-beta straddle returns corresponding to strike prices immediately above and below each day's opening level of the S&P 100. Then we cumulate these daily returns for each month. Because our straddle returns are highly sensitive to innovations in volatility, their excess returns should capture any ability of volatility risk to account for cross-sectional variation in excess returns. For each size portfolio i , we conduct a regression of the following form:

$$r_{it} - r_f = \alpha_i + \beta_{im}(r_{mt} - r_f) + \beta_{iv}(r_{vt} - r_f) + \varepsilon_{it}, \quad (22)$$

where r_{it} reflects the return on portfolio i in period t , r_f is the risk-free rate of return, r_{mt} is the period t market return, and r_{vt} represents the period t zero-beta straddle return. The regression results are presented in Table VIII.

As we can see, the betas on the volatility factor are significant across all 10 portfolios. Most interesting is that the coefficients are closely related to portfolio size. The magnitudes range begin with -0.0328 on portfolio 1 and increase steadily to 0.0059 for the largest size portfolio. This indicates that a firm's size may be closely related to its exposure to volatility risk. Small firms, it appears, experience far worse returns when volatility increases than larger firms. In fact, firms in the largest size portfolio appear to hedge their investors against innovations to volatility. This result, if sustained, offers a potential explanation for size-related return anomalies. Large firms, in hedging their investors against increases in market risk, earn lower equilibrium returns than those justified by their CAPM betas.

IV. Conclusion

Although the pricing of option contracts is the subject of considerable attention from researchers, option returns have been largely neglected in the literature. Standard asset-pricing theory claims that option risks should be priced no differently than those of other assets in the economy—according to the degree of systematic risk they require their holders to bear. Hence, option markets offer an excellent setting in which to test basic implications of asset-pricing theory.

In this paper, we offer the first detailed examination of long-run option returns in the context of implications set forth by asset-pricing theory. We find that option returns largely conform to most asset-pricing implications. However, considering their levels of systematic risk, we find considerable evidence that both call and put contracts earn exceedingly low returns. A strategy of buying zero-beta straddles has an average return of around -3 percent per week. This finding is robust to a variety of checks, including allowing for measurement errors, altering the sample period and frequency, removing the crash-insurance component of the straddle premium, and accounting for transaction costs.

Our results strongly suggest that something besides market risk is important for pricing the risk associated with option contracts. Because zero-beta straddle returns are basically determined by innovations in market volatility, our results imply that systematic stochastic volatility may be an important factor for pricing assets. This interpretation is supported by our findings that straddle returns are also negative in other options markets and are somewhat related to the correlation between their implied volatilities and those of the S&P 100 index. We also offer some preliminary evidence that a volatility factor helps explain a significant component of the time-series variation in portfolio returns. Intriguingly, the returns of large firms are positively related to innovations in volatility, offering a possible explanation for size-based returns anomalies.

It would be natural to extend the results in this paper by fitting the data with models that incorporate systematic stochastic volatility like Heston (1993). However, there is already some evidence that Heston's model does not fit the

data (see Bates (1997) and Chernov and Ghysels (1998)). Also, it is important to note that simply incorporating crash risk will not be sufficient to explain the nature of the option returns documented here. Another natural extension of our results might be an investigation of whether volatility risk is priced in the returns of other assets. If volatility risk were found to be priced in a broader set of assets, it might shed light on existing cross-sectional asset-pricing puzzles.

REFERENCES

- Bates, David S., 1997, Post-'87 crash fears in S&P 500 futures options, Working paper 5894, NBER.
- Black, Fischer, and Myron Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–659.
- Britten-Jones, Mark, 1999, The sampling error in estimates of mean-variance efficient portfolio weights, *Journal of Finance* 54, 655–671.
- Buraschi, Andrea, and Jens Jackwerth, 1999, Is volatility risk priced in the option market? Working paper, London Business School.
- Canina, Linda, and Stephen Figlewski, 1993, The information content of implied volatility, *Review of Financial Studies* 6, 659–681.
- Chernov, Mikhail, and Eric Ghysels, 1998, What data should be used to price options? Working paper, Pennsylvania State University.
- Christensen, B. J., and N. R. Prabhala, 1998, The relation between implied and realized volatility, *Journal of Financial Economics* 50, 125–150.
- Dybvig, Phillip H., 1981, A practical framework for capital budgeting of projects having uncertain returns, Working paper, Washington University.
- Gibbons, Michael R., Stephen A. Ross, and Jay Shanken, 1989, A test of the efficiency of a given portfolio, *Econometrica* 57, 1121–1152.
- Heston, Steven 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.
- Jackwerth, Jens, 2000, Recovering risk aversion from option prices and realized returns, *Review of Financial Studies* 13, 433–451.
- Jackwerth, Jens, and Mark Rubinstein, 1996, Recovering probability distributions from contemporary security prices, *Journal of Finance* 51, 1611–1631.
- Merton, Robert C., 1971, Optimum consumption and portfolio rules in a continuous-time model, *Journal of Economic Theory* 3, 373–413.
- Merton, Robert C., Myron S. Scholes, and Matthew L. Gladstein, 1978, The returns and risk of alternative call option portfolio investment strategies, *Journal of Business* 51, 183–242.
- Merton, Robert C., Myron S. Scholes, and Matthew L. Gladstein, 1982, The returns and risk of alternative put option portfolio investment strategies, *Journal of Business* 55, 1–55.
- Sheikh, Aamir, and Ehud I. Ronn, 1994, A characterization of the daily and intraday behavior of returns on options, *Journal of Finance* 49, 557–579.