The price of market volatility risk

Jefferson Duarte* and Christopher S. Jones†

October 28, 2007

Abstract

We analyze the volatility risk premium by applying a modified two-pass Fama-MacBeth procedure to the returns of a large cross section of the returns of options on individual equities. Our results provide strong evidence of a volatility risk premium that is increasing in the level of overall market volatility. This risk premium provides compensation for risk stemming both from the characteristics of the option contract and the riskiness of the underlying equity. We also show with a large scale Monte Carlo simulation that measurement error in option prices and violations of arbitrage bounds induce highly economically significant biases in the mean returns of options. In fact, our simulation results demonstrate that biases can be up to several percentage points per day. These large biases can lead researchers to faulty conclusions with respect to both the magnitude of the volatility risk premium and the sign of expected option returns.

*University of Washington. E-mail: jduarte@u.washington.edu
†University of Southern California. E-mail: christopher.jones@marshall.usc.edu

1 Introduction

There is no doubt that stock market volatility changes over time, but whether or not volatility represents a priced risk factor remains less certain. Theoretically, a stochastic volatility factor seems a prime candidate for a nonzero risk premium because, in the framework of Merton's (1973) ICAPM, it represents a variable that drives the investment opportunity set. Empirically, support for a nonzero volatility risk premium is varied, with some studies finding large values and some finding none. The primary goal of this paper is to address this question using an empirical method and a data set that together should offer a much clearer answer than has previous work.

Determining whether volatility is priced has important consequences for option pricing, interpreting implied volatilities, and understanding investor preferences. Stochastic volatility models are the norm in the option pricing literature, and numerous papers have demonstrated that a nonzero price of volatility risk can improve model fit substantially. A volatility risk premium also means that implied volatilities, whether they are from the Black-Scholes (1973) model or the model-free approach of Britten-Jones and Neuberger (2000), cannot be interpreted as unbiased forecasts of future realized volatility. An assessment of the size of the volatility risk premium is also useful in understanding the preferences of market participants and in building better models of investor behavior. Finally, knowing whether a volatility risk premium exists has important implications for understanding the purpose of options markets, namely the extent to which these can be seen as markets for volatility risk.

In this paper we apply traditional two-pass Fama-MacBeth regressions to the returns of a large cross section of options on individual equities. Analyzing equity options rather than index options, as virtually all prior work has done, introduces an additional source of variation in volatility factor betas. Specifically, in any analysis of index options, all variation in volatility betas must arise from variation in the characteristics of the different option contracts, namely time-to-maturity, moneyness, and whether the option is a call or a put. Because volatility factor betas are likely correlated with a number of other variables that could be relevant in determining risk premia, it is possible that volatility betas are merely a proxy for some other source of expected return that is also related to maturity and moneyness.

Looking at options on individual equities allows for an additional source of variation in volatility betas that is independent from the characteristics of the option contract. Specifically, using multiple underlying securities allow us to sort firms into groups that appear to have a greater degree of dependence between their own return volatilities and the volatility of the market, and groups that do not. We can use this sorting procedure to introduce additional dispersion into the betas of option portfolios. We can also use the procedure to isolate the variation that is due solely to firm-level rather than contract-level variation.

We believe that the use of the traditional Fama-MacBeth approach posesses some significant advantages relative to recent work that emphasizes fully specified parametric and usually continuous-time models. The most important of these advantages is related to the highly rigid structure placed on risk premia by most parametric option pricing models. In a pure stochastic volatility model, for example, any abnormal return on a strategy of buying zero delta straddles must be attributed to a nonzero price ov volatility risk – there are simply no other channels within the model to generate this result. Furthermore, parametric models put substantial structure on the relative riskiness of different option contracts. Even if the prices of risk are correct, misspecification in risk exposures will generate misleading conclusions about the importance of risk premia in explaining option pricing anomalies. Of course, good empirical work typically examines these assumption in one manner or another, and we are not suggesting that the estimation of continuous-time option pricing models be discontinued, only that a less parametric approach might provide a valuable complement to the existing literature.

Unfortunately, we find that traditional "off the shelf" methods cannot be applied in the setting in which we are interested. Option markets are characterized by large bid-ask spreads, and the noise that these spreads inject into observations of fair market value causes several significant problems. First, noisy prices may not conform to arbitrage bounds, usually because the observed option price is below its intrinsic value. This makes it impossible to solve for an implied volatility, for example, which is typically the first step in computing the delta used to hedge the option. Prior work has usually dealt with this issue by throwing out observations that do not satisfy these bounds, but because most of the discarded observations are of option prices that are too low rather than too high, this method introduces a censoring bias. Secondly, as Blume and Stambaugh (1982) first pointed out, noisy prices induce an upward bias in average returns.

Our empirical work is guided by an extensive simulation study that focuses on what modifications of the standard two-pass approach are necessary to achieve the correct test size and a high degree of power. Through this process, we identify a successful alternative to the traditional approach of throwing out observations that do not satisfy arbitrage bounds. We also implement a new method for bias adjustment that allows us to reduce, if not eliminate completely, the bias identified by Blume and Stambaugh. Our simulation results demonstrate that both biases can be severe, with each causing biases in average returns of up to several percentage points *per day.* Although these simulation results might appear overly dramatic, we have found virtually identical results when we look at actual data. At the same time, our analysis shows that other common concerns about applying traditional methods to option returns appear to be overrated – there appears to be little worry, for instance, that nonlinearity in the return-factor relation is driving our results.

When we apply our modified two-pass approach to the data, our results strongly support the view that stochastic volatility represents a priced risk factor, but only conditionally as the unconditional mean cannot reliably be distinguished from zero. In particular, we find strong evidence that the volatility risk premium varies positively with the level of implied volatilities from S&P 500 index options. This result is robust to a variety of different portfolio construction methods and variations in our econometric approach. In addition, we find strong evidence that these risk premia are not merely proxies for other variables that are similarly correlated with option contract characteristics such as maturity and strike price, as the dispersion in volatility betas that is driven by differences in underlying stocks is priced similarly to the dispersion coming from option contract type.

In a brief application of our framework to equity index options, we demonstrate that the bias caused by bid-ask bounce could be sufficient to reverse the observed relation between strike prices and average call option returns. Like Coval and Shumway (2001), we find a positive relation between strike price and average observed returns, a result that they note is consistent with the the higher CAPM betas of out-of-the-money (OTM) call options. Our regression results indicate, however, that this effect could be entirely driven by the bias arising from higher bid-ask spreads among OTM options. After correcting for this bias, we find that the result reverses, suggesting that the CAPM does not provide an adequate explanation for the expected returns of these assets. These findings lead us to believe that similar biases might help explain some of the other puzzles that have been identified in the option pricing literature.

The prior literature on volatility risk premia is large, but almost all of it focuses on equity index options. Early tests of risk premia focused on comparing future realized volatility with current Black-Scholes implied volatilities, usually in a regression framework. Papers such as Fleming, Ostdiek, and Whaley (1995), Jackwerth and Rubinstein (1996), and Christensen and Prabhala (1998), show that realized volatilities tend to be substantially lower than implied volatilities. Subsequent work on the estimation of structural models, most based on the square root specification of Heston (1993), offer complementary evidence that confirmed that the results are not artifacts of the misuse of the Black-Scholes formula. In particular, papers by Bates (2000), Benzoni (2002), Chernov and Ghysels (2000), Jones (2003), and Pan (2002) all find large negative volatility risk premium. A slightly dissenting voice is Broadie, Chernov, and Johannes (2005), who find no evidence for a risk premium on *diffusive* volatility risk but possibly a risk premium on volatility jumps. Finally, the presence of a risk premium is suggested by the returns-based analyses of Coval and Shumway (2001) and Jones (2003), who find that expected returns on short delta-hedged options are often on the order of half a percent per day, indicating the presence of a residual risk factor with a large risk premium.

Work analyzing individual equity options is relatively sparse. Bakshi and Kapadia (2003) present some evidence, but the conclusions are weak most likely because their is only five years long and analyzes only 25 firms. Carr and Wu (2003) use a similarly limited, though slightly longer sample and find slightly stronger evidence for a volatility risk premium, though they find it to be weaker in equity options than in index options. On the other hand, Driessen, Maenhout, and Vilkov (2006) compare average model-free implied variances to average realized variances and find insignificant differences, which they cite as evidence that the volatility risk premium does not in fact exist.

Finally, both Chen (2002) and Ang, Hodrick, Xing, and Zhang (2006) analyze whether volatility is a priced factor in the cross-section of equity returns. While Chen's (2002) GARCH-based approach results in inconclusive results, Ang et al. use a factor constructed from index implied volatilities and find a significantly negative risk premium.

Our paper proceeds as follows. In Section 2 we discuss the basic modeling framework used throughout the paper. Section 3 introduces the econometric methods we rely on, while Section 4 includes a description of our data set. Section 5 describes our Monte Carlo study and presents its results. Section 6 contains our main empirical findings, and Section 7 concludes.

2 Theory

In this section we motivate our approach using a standard continuous-time stochastic volatility model. We discuss the risk exposures faced both by naked and hedged option buyers, distinguishing standard delta hedging from a more theoretically accurate approach for eliminating underlying risk using a measure we call "total delta." Risk premia are first introduced mechanically as risk prices corresponding to the Brownian motions driving each firm, but we then link these risk prices to exposure to market-level risk.

2.1 A stochastic volatility model of equity prices

We start with the following general stochastic volatility model:

$$\begin{aligned} \frac{dS_t}{S_t} &= \mu_t dt + \sigma_t dB_{1,t} \\ d\sigma_t &= \theta_t dt + \omega_t \rho dB_{1,t} + \omega_t \sqrt{1-(\rho)^2} dB_{2,t} \\ \text{corr}\left(\frac{dS_t}{S_t}, d\sigma_t\right) &= \rho, \end{aligned} \tag{1}$$

where S_t is the price of the stock at time t, σ_t is its volatility, and ρ is the instantaneous correlation between the stock price and its volatility. The Brownian motions $B_{1,t}$ and $B_{2,t}$ are standard Brownian motions under the physical probability measure.

The price of a derivative on the underlying stock is represented by $f(S_t, \sigma_t, t)$. Under standard assumptions, Appendix A shows that the instantaneous excess return of this derivative is

$$\begin{aligned}\frac{df_t}{f_t} - r_t dt &= \left(\frac{S_t}{f_t}\frac{\partial f}{\partial S} + \frac{\omega_t \rho}{\sigma_t}\frac{1}{f_t}\frac{\partial f}{\partial \sigma}\right)\left(\frac{dS_t}{S_t} - r_t dt + q_t dt\right) \\ &\quad + \omega_t\sqrt{1-\rho^2}\frac{1}{f_t}\frac{\partial f}{\partial \sigma}\left(\lambda_{2,t}dt + dB_{2,t}\right), \end{aligned} \tag{2}$$

where q_t is the instantaneous dividend yield of stock i and r_t is the instantaneous risk-free rate of interest. Both q_t and r_t are possibly time-varying but deterministic. The expected excess return is therefore given by the following "two-beta" representation:

$$\mathrm{E}\left[\frac{df_t}{f_t}\right] - r_t dt = \left(\frac{S_t}{f_t}\frac{\partial f}{\partial S} + \frac{\omega_t \rho}{\sigma_t}\frac{1}{f_t}\frac{\partial f}{\partial \sigma}\right)\lambda_{1,t}\sigma_t dt + \frac{1}{f_t}\frac{\partial f}{\partial \sigma}\lambda_{2,t}\omega_t\sqrt{1-\rho^2}dt \tag{3}$$

Thus, the parameter $\lambda_{1,t}$ is the price of stock risk and $\lambda_{2,t}$ is the price of *orthogonal* volatility risk. The price of stock risk is also by definition analogous to an instantaneous Sharpe ratio, that is $\lambda_{1,t} = (\mu_t - r_t + q_t)/\sigma_t$.

The expected excess return of the derivative in Equation (3) is composed of two parts, one related to the exposure to stock price risk ($\lambda_{1,t}$), the other related to the exposure to the volatility risk ($\lambda_{2,t}$). The exposure of derivatives to stock returns is also composed of two parts, one related to the direct relation between stock prices and derivative prices ($\partial f/\partial S$), and the other related to the correlation between volatility and stock returns (ρ).

The expected excess returns of both call and put options in the model above can be either positive or negative. Under the assumption of deterministic volatility, Coval and Shumway (2001) show that the expected excess returns of calls is positive while the expected excess returns of puts is negative. Mathematically, this result of Coval and Shumway (2001) is a direct consequence of Equation (3), which under the assumption of deterministic volatility is reduced to:

$$E\left[\frac{df}{f}\right] - r dt = \sigma\frac{S}{f}\frac{\partial f}{\partial S}\lambda_{1,t}dt \tag{4}$$

The intuition for this result is that if volatility is deterministic, the expected excess return of a derivative results from its exposure to stock price risk and the expected excess return of the underlying stock. Because call (put) options are positively (negatively) exposed to stock prices and stock expected excess returns are generally positive, the expected excess return of calls is positive and the expected excess return of puts is negative. Under the assumption of stochastic volatility, on the other hand, this conclusion may not hold, because both calls and puts are also exposed to volatility risk. Consequently, the expected excess return of calls and puts depends on how their exposure to the volatility risk balances with their exposure to the stock price risk. For instance, assume that the correlation between stock returns and volatility changes is zero ($\rho = 0$), the price of volatility risk is negative ($\lambda_{1,t} < 0$) and the price of stock risk is positive ($\lambda_{2,t} > 0$). In this case, the expected excess return of call options could be negative given that the call option vega ($\nu = \partial f / \partial \sigma$) must be positive.

2.2 Regular delta versus "total" delta

To empirically distinguish between the effects of stock-price risk and volatility, we compute the return of options hedged against the exposure to stock price risk. The hedged portfolios are composed by a long position in a single option contract and a position on n stocks financed at the risk-free rate (r_t), where n is given by:

$$n = -(\frac{\partial f}{\partial S} + \frac{\partial f}{\partial \sigma}\frac{\omega\rho}{\sigma S}) = -(\Delta + \nu\frac{\omega\rho}{\sigma S}) \tag{5}$$

Appendix B shows that hedging one option contract with n shares of stock provides a complete hedge against any exposure to the underlying stock price. Note that the number of shares reduces to minus the delta of the options ($\Delta = \partial f / \partial S$) in the event that stock returns are uncorrelated with changes in volatility. On the other hand, when volatility and stock returns are correlated it is necessary to change n to account for the effects that common movements on stock prices and volatilities may have on options prices. We call the hedging procedure with n shares of a *total-delta hedge* to distinguish from the usual delta-hedging procedure in which the number of shares used to hedge a derivative is set equal to $-\partial f / \partial S$.

The expected excess return of the hedged portfolio does not depend on the price of stock price risk ($\lambda_{1,t}$). As shown in Appendix B, the excess return of a portfolio of a option delta-hedged option is:

$$\frac{dH}{H} - rdt = \omega\sqrt{1-\rho^2}\lambda_{2,t}\frac{1}{f}\frac{\partial f}{\partial \sigma}dt + \frac{1}{f}\frac{\partial f}{\partial \sigma}\omega\sqrt{1-\rho^2}dB_{2,t}, \tag{6}$$

where H represents the value of a portfolio composed of an option hedged with n stocks. Note that the mean return of the total-delta hedged option depends on the price of volatility risk ($\lambda_{2,t}$). This mean return is the compensation required by investors to hold the risk related to volatility movements not common to stock price movements.

2.3 Sources of risk premia

Assume that the market portfolio follows the following dynamics under the physical probability measure:

$$\frac{dS_t^m}{S_t^m} = \mu_t^m dt + \sigma_t^m dB_{1,t}^m \tag{7}$$

$$
\begin{aligned}
d\sigma_t^m &= \theta_t^m dt + \omega_t^m \rho^m dB_{1,t}^m + \omega_t^m \sqrt{1-(\rho^m)^2} dB_{2,t}^m \\
\text{corr}\left(\frac{dS_t^m}{S_t^m}, d\sigma_t^m\right) &= \rho^m
\end{aligned}
$$

where $B_{1,t}^m$, $B_{2,t}^m$ are uncorrelated Brownian motions. The Brownian motions affecting the individual stock return and volatility can be rewritten as:

$$
\begin{aligned}
dB_{1,t}^i &= \xi_{11}^i dB_{1,t}^m + \xi_{12}^i dB_{2,t}^m + dZ_{1,t}^i \\
dB_{2,t}^i &= \xi_{21}^i dB_{1,t}^m + \xi_{22}^i dB_{2,t}^m + dZ_{2,t}^i,
\end{aligned} \tag{8}
$$

where $B_{1,t}^i$ and $B_{1,t}^i$ are the same Brownian motions in Equation (1). (The superscript i is added to avoid confusion with the market-related Brownian motions.) The Brownian motions $Z_{1,t}^i$ and $Z_{2,t}^i$ have zero drift and quadratic variation equal to $1-(\xi_{11}^i)^2-(\xi_{12}^i)^2$ and $1-(\xi_{21}^i)^2-(\xi_{22}^i)^2$, respectively. Furthermore, both $Z_{1,t}^i$ and $Z_{2,t}^i$ are uncorrelated with $B_{1,t}^m$ and $B_{2,t}^m$. Thus, ξ_{11}^i, ξ_{12}^i, ξ_{21}^i, and ξ_{22}^i represent the correlations between the Brownian motions driving the stock process and those driving the market process. These correlations are assumed to be constant. In the decomposition above the individual stock returns and volatilities are decomposed into components that are related to market returns and to changes in market volatility and into components that are idiosyncratic. Assuming that the price of idiosyncratic risk is zero, we conclude that the prices of stock and volatility risks are:

$$
\begin{aligned}
\lambda_{1,t}^i &= \xi_{11}^i \lambda_{1,t}^m + \xi_{12}^i \lambda_{2,t}^m \qquad (9) \\
\lambda_{2,t}^i &= \xi_{21}^i \lambda_{1,t}^m + \xi_{22}^i \lambda_{2,t}^m
\end{aligned}
$$

where the price of market risk is $\lambda_{1,t}^m = (\mu_t^m - r_t + q_t^m)/\sigma_t^m$, where q_t^m is the instantaneous dividend yield paid by the market portfolio.

The price of stock risk ($\lambda_{1,t}^i$) in the Equation (9) is analogous to the standard beta representation of expected returns. To see this, note that multiplying $\lambda_{1,t}^i$ in Equation (9) by σ_t^i, gives:

$$
\begin{aligned}
(\mu_t^i - r_t + q_t^i) &= \beta_{m,t}^i(\mu_t^m - r_t + q_t^m) + \beta_{Vol,t}^i(\omega_t^m \sqrt{1-(\rho^m)^2}\lambda_{2,t}^m) \\
\beta_{m,t}^{i,1} &= \frac{\sigma_t^i \xi_{11}^i}{\sigma_t^m} \\
\beta_{Vol,t}^{i,1} &= \frac{\sigma_t^i \xi_{12}^i}{\omega_t^m \sqrt{1-(\rho^m)^2}}
\end{aligned} \tag{10}
$$

The first term on the right hand side of the first equation above is analogous to the standard $CAPM$ relation between the expected excess return of a stock and the expected excess return of the market, as $\beta_{m,t}^i$ is the stock beta in this context. The second term is the premium required by investors to hold a stock whose returns covary with the volatility of the entire market. The term $\beta^{i,1}_{Vol,t}$ is the loading of this stock to the market volatility risk and $(\omega^m_t\sqrt{1-(\rho^m)^2}\lambda^m_{2,t})$ is the premium required by investors for holding the market volatility risk.

The price of volatility risk $(\lambda^i_{2,t})$ in the Equation (9) also implies a beta representation that is

$$\begin{aligned}
\omega^i_t\sqrt{1-(\rho^i)^2}\lambda^i_{2,t} &= \beta^{i,2}_{m,t}(\mu^m_t - r_t + q^m_t) + \beta^{i,2}_{Vol,t}(\omega^m_t\sqrt{1-(\rho^m)^2}\lambda^m_{2,t}) \\
\beta^{i,2}_{m,t} &= \frac{\omega^i_t\sqrt{1-(\rho^i)^2}\xi^i_{21}}{\sigma^m_t} \\
\beta^{i,2}_{Vol,t} &= \frac{\omega^i_t\sqrt{1-(\rho^i)^2}\xi^i_{22}}{\omega^m_t\sqrt{1-(\rho^m)^2}}
\end{aligned} \tag{11}$$

In this representation $\beta^{i,2}_{m,t}$ is the beta of a total-delta hedged portfolio with respect to the market and $\beta^{i,2}_{Vol,t}$ is the beta of the individual stock total-delta hedged options portfolio with respect to the market volatility factor.

Equation (11) is the center of the Fama-MacBeth procedure in Section 3.3. In the first step of our Fama-MacBeth procedure, we estimate $\beta^{i,2}_{m,t}$ from the relation between the returns of a portfolio of total-delta-hedged options on the individual stock i and the market portfolio returns , we estimate $\beta^{i,2}_{Vol,t}$ from the relation between the returns of a portfolio of total-delta-hedged options on the individual stock i and the returns of a portfolio of total-delta-hedged options written on the market portfolio. In the second step of our Fama-MacBeth procedure, the risk premium of volatility risk $(\omega^m_t\sqrt{1-(\rho^m)^2}\lambda^m_{2,t})$ inferred from the cross-section of options is estimated.

3 Econometric approach

We begin this section by providing an empirical strategy for implementing the "total delta" hedge described in the previous section. We then discuss the process we use for forming portfolios that maximize dispersion in volatility betas. Finally, we discuss the implementation of the Fama-Macbeth procedure and how we address three potential sources of bias.

3.1 Computing hedge ratios

In general, our hedging strategies rely on implied volatilities and "Greeks" from the Black-Scholes model. This approach, while approximate, is standard practice in industry and has been shown in academic research to be quite accurate. Hull and Suo (2002), in particular, go so far as to claim that "the pricing and hedging of a new or an existing plain vanilla instrument is largely model independent." An important qualifier to this statement is that hedge ratio inputs, such as deltas, must be computed using the implied volatilities of the option being hedged. We follow this practice below.

Most of the inputs required to compute hedge ratios can be inferred from options prices using the Black-Scholes formula. For instance, to hedge a long position in one option, we must take a position of n shares financed at the risk-free rate, where the hedge ratio n is given by equation (5). Consequently, the hedging procedure requires knowledge of deltas $(\partial f/\partial S)$, vegas $(\partial f/\partial \sigma)$, and the term $\omega\rho/\sigma$ We take the delta and the vega to be equal to their Black-Scholes values (computed using the option's own implied volatility). The last term, however, cannot be directly observed and must instead be estimated using time series regression.

To estimate $\omega\rho/\sigma$, we assume that ω is a constant for each firm. This is an implication of Heston's (1993) model of stochastic volatility. We also assume that ρ is constant, and we take σ to equal the implied volatility on the option. Our task is then to estimate the constant $\omega\rho$. From Equation (47) in Appendix A, we can write the return on hedged option, where the hedge ratio is computed using the standard definition of delta, as

$$\frac{df}{f} - \frac{S\sigma}{f}\frac{\partial f}{\partial S}dB_{1,t} = \text{drift} + \frac{\omega\rho}{f}\frac{\partial f}{\partial \sigma}dB_{1,t} + \frac{\omega\sqrt{1-\rho^2}}{f}\frac{\partial f}{\partial \sigma}dB_{2,t} \tag{12}$$

Since $dS_t/S_t = \mu_t^i dt + \sigma_t dB_{1,t}^i$, we can rewrite this equation as

$$\frac{df}{f} - \frac{S}{f}\frac{\partial f}{\partial S}\frac{dS_t}{S_t} = \text{drift} + \frac{\omega\rho}{\sigma f}\frac{\partial f}{\partial \sigma}\frac{dS_t}{S_t} + \frac{\omega\sqrt{1-\rho^2}}{f}\frac{\partial f}{\partial \sigma}dB_{2,t} \tag{13}$$

This implies that we can estimate the quantity $\omega\rho$ as the slope coefficient from the regression

$$y_t = a + b\frac{dS_t}{S_t} + \epsilon_t, \tag{14}$$

where y_t is the average, across all options in a given day for a given stock, of the quantity

$$\left(\frac{df}{f} - \frac{S}{f}\frac{\partial f}{\partial S}\frac{dS_t}{S_t}\right) / \left(\frac{1}{f\sigma}\frac{\partial f}{\partial \sigma}\right). \tag{15}$$

Note that the expression above is written in terms of infinitesimal increments, while our data are observed discretely. We therefore use a discrete approximation to the continuous time model to compute this quantity. That is, the quantity above is estimated in a given day t with the approximation

$$\left(\frac{f_t - f_{t-1}}{f_{t-1}} - \frac{S_{t-1}}{f_{t-1}}\Delta_{t-1}\frac{S_t - S_{t-1}}{S_{t-1}}\right) / \left(\frac{1}{f_{t-1}\sigma_{t-1}}\nu_{t-1}\right), \tag{16}$$

where σ_{t-1} is the at-the-money volatility of the stock.

3.2 Estimating risk factor sensitivities

Given a strategy to estimate all the variables required to calculate n, we can now compute the excess return on "total delta-heged" options. This return should, if constructed accurately, be independent of any movements in the underlying security, even those that arise through a nonzero correlation between the stock price and volatility processes. It therefore isolates the component of option returns that is unambiguously related to volatility, and it is the primary subject of our analysis.

Letting H^{ij} denote the value of a hedged position in option contract j on security i, we can use the decomposition of B_2 in equation (8) and the form of the price of risk in equation (9) to rewrite (6) as

$$\begin{aligned} \frac{dH^{ij}}{H^{ij}} - rdt &= \frac{1}{f^{ij}} \frac{\partial f^{ij}}{\partial \sigma^i} \omega^i \sqrt{1-(\rho^i)^2} \times \\ & \left[\xi^i_{21} \left(\lambda^m_1 dt + dB^m_{1,t} \right) + \xi^i_{22} \left(\lambda^m_2 dt + dB^m_{2,t} \right) + dZ^i_{2,t} \right] \end{aligned} \tag{17}$$

This representation is useful because it isolates the two possible sources of systematic risk in hedged option returns. One (ξ^i_{21}) is the sensitivity to a component of stock market returns that is independent of the returns on equity i. The other (ξ^i_{22}) is the sensitivity of equity i's volatility process to market volatility.

Note, however, that these sensitivities are only transferred to hedged option returns to the extent the option is sensitive to volatility movements in general. This sensitivity, which we call the *normalized vega*, is equal to $(1/f^{ij} \times \partial f^{ij}/\partial \sigma^i)$. If we divide both sides of the equation by this amount, we get a term that should be identical across all options on the same underlying stock:

$$\begin{aligned} \Pi^{ij} &\equiv \left(\frac{dH^{ij}}{H^{ij}} - rdt \right) / \left(\frac{1}{f^{ij}} \frac{\partial f^{ij}}{\partial \sigma^i} \right) \\ &= \omega^i \sqrt{1-(\rho^i)^2} \left[\xi^i_{21} \left(\lambda^m_1 dt + dB^m_{1,t} \right) + \xi^i_{22} \left(\lambda^m_2 dt + dB^m_{2,t} \right) + dZ^i_{2,t} \right] \end{aligned} \tag{18}$$

Because the Π^{ij} should be identical for all options j on a given stock i, for a given date, we will look at averages across options on the same stock, which we denote as Π^i. Note that Π^i can be interpreted as a return on a zero investment strategy involving a number of hedged option positions. It is also the excess capital gain resulting from holding a portfolio with vega ($\partial f/\partial \sigma$) equal to one. To see this note that by construction our hedged portfolios have the same value as the option prices themselves (i.e. $H = f$), and hence an option with vega equal to one would have capital gains over time dt equal to Π^i. Because of this interpretation we call this average the *one-vega P&L*.

The same derivations also apply to the options on the S&P 500 index, which we view as a proxy of the market portfolio, except that by definition of B_2^m we have

$$\Pi^{mj} \equiv \left(\frac{dH^{mj}}{H^{mj}} - rdt \right) / \left(\frac{1}{f^{mj}} \frac{\partial f^{mj}}{\partial \sigma^m} \right) = \omega^m \sqrt{1-(\rho^m)^2} \left(\lambda_2^m dt + dB_{2,t}^m \right) \tag{19}$$

Define Π^m as the average of these values across all options for a given date. We call this average the *market one-vega P&L*. Equation (19) implies that the average value Π^{mj} is an estimate of the volatility risk premium $\omega^m \sqrt{1-(\rho^m)^2} \lambda_2^m dt$.

A consequence of the definitions of the market one-vega P&L and of the stock one-vega P&L is that we can use a time-series regression of a given stock one-vega on the market return and on the market one-vega to estimate $\beta_{m,t}^{i,2}$ and $\beta_{m,t}^{i,2}$ in Equation (11). To see this, note that we can write the stock one-vega as

$$\Pi^i = \beta_{m,t}^{i,2} \left(\frac{dS^m}{S^m} - rdt + qdt \right) + \beta_{Vol,t}^{i,2} \Pi^m + \omega^i \sqrt{1-(\rho^i)^2} dZ_{2,t}^i \tag{20}$$

In this representation $\beta_{m,t}^{i,2}$ is the beta of a delta-hedged portfolio with respect to the market and $\beta_{Vol,t}^{i,2}$ is the beta of the individual stock delta-hedged options portfolio with respect to the market volatility factor.

A concern is that neither beta is likely to be constant. For instance, $\beta_{m,t}^{i,2}$ will vary inversely with market volatility unless the quantity $\omega^i \sqrt{1-(\rho^i)^2} \xi_{21}^i$ happens to be proportional to it. The assumption of the constancy of $\beta_{Vol,t}^{i,2}$ is also unsupported, although it will be the case in the parameterization of the Heston (1993) model that we consider the Monte Carlo experiments in Section 5. In general, we view the estimation of these betas via time series regression as simply an approximate way to rank stocks according to the degree of their systematic volatility risk.

3.3 Fama-MacBeth regression

We use Fama-MacBeth regressions to understand the effect of volatility risk on expected returns of options. We are interested in understanding whether the magnitude of the price of volatility risk implied by individual stock options and in understanding whether the premium for volatility risk in individual stock options stems from the relation between individual stock volatilities and overall market volatility. To do so, we examine portfolios of options formed on the basis of maturity, moneyness, and estimates of $\beta_{Vol,t}^{i,2}$. We sort on $\beta_{Vol,t}^{i,2}$ rather than $\beta_{m,t}^{i,2}$ because we expect the former to contribute more to the variation in expected returns given the large risk premia that appear to exist in equity index option prices. We estimate $\beta_{Vol,t}^{i,2}$ using the time series regression described in Section 3.2 with data from the six months preceding the portfolio formation date. Regressions are re-estimated and portfolios are rebalanced daily.

Portfolios are equally weighted across options, and portfolio returns are computed as averages of hedged option excess returns, $dH^{ij}/H^{ij} - r_t dt$. From Section 3.2, we can see that this return can be rewritten as

$$\frac{1}{f^{ij}}\frac{\partial f^{ij}}{\partial \sigma^i}\left[\beta^{i,2}_{m,t}\left(\frac{dS^m}{S^m} - rdt\right) + \beta^{i,2}_{Vol,t}\Pi^m + \omega^i\sqrt{1-(\rho^i)^2}dZ^i_{2,t}\right] \tag{21}$$

The portfolio return can therefore be written as

$$\frac{1}{N}\sum_{ij}\left[\frac{dH^{ij}}{H^{ij}} - rdt\right] = \frac{1}{N}\sum_{ij}\left[\frac{1}{f^{ij}}\frac{\partial f^{ij}}{\partial \sigma^i}\beta^{i,2}_{m,t}\right]\left(\frac{dS^m}{S^m} - rdt\right) \tag{22}$$

$$+\frac{1}{N}\sum_{ij}\left[\frac{1}{f^{ij}}\frac{\partial f^{ij}}{\partial \sigma^i}\beta^{i,2}_{Vol,t}\right]\Pi^m + \frac{1}{N}\sum_{ij}\left[\frac{1}{f^{ij}}\frac{\partial f^{ij}}{\partial \sigma^i}\omega^i\sqrt{1-(\rho^i)^2}dZ^i_{2,t}\right]$$

where the last term is an error that is uncorrelated with aggregate risk factors.

As in the single-firm analysis above, the expectation of this return is linear in two betas, namely

$$\frac{1}{N}\sum_{ij}\left[\frac{1}{f^{ij}}\frac{\partial f^{ij}}{\partial \sigma^i}\beta^{i,2}_{m,t}\right] \quad \text{and} \quad \frac{1}{N}\sum_{ij}\left[\frac{1}{f^{ij}}\frac{\partial f^{ij}}{\partial \sigma^i}\beta^{i,2}_{Vol,t}\right] \tag{23}$$

Also similar to the single-firm analysis is the fact that neither beta is likely to be constant. This is not only due to the variation in firm-level $\beta^{i,2}_{m,t}$ and $\beta^{i,2}_{Vol,t}$, but also because time variation in the composition of each of the portfolios will cause some variation in normalized vegas. It seems likely, however, that some of this variation will be eliminated by averaging across the large number of contracts in each portfolio.

Our approach is therefore to estimate these portfolio betas from the full sample of triple-sorted portfolio returns and to use this single set of betas in a second stage Fama-MacBeth regression.

3.4 Implementation issues

3.4.1 Discretization issues

As noted above, the regressions are written in terms of infinitesimal increments, while our data are observed discretely. We are therefore required to use a discrete approximation to the continuous time model, and for tractability the approximation we use is linear, i.e. $dX_t \approx X_t - X_{t-1}$ and $dX_t/X_t \approx X_t/X_{t-1} - 1$. A natural concern is that the linear approximation is inadequate. Branger and Schlag (2004), for instance, argue that even in a Black-Scholes economy the expected return on a delta-hedged portfolio will be nonzero if the hedge is rebalanced with moderate frequency.

This bias arises from the fact that the pricing kernel in the Black-Scholes (or any other option pricing) model is nonlinear. As a simple illustration, suppose the pricing kernel were equal to $a + bR^m + c(R^m)^2$. A delta-hedged return should be orthogonal to the linear term, but the convexity of option payouts makes it likely to be correlated with $(R^m)^2$. Even in a Black-Scholes world, this should result in a nonzero expected excess return.

We address this issue first by exclusively using daily data, for which the nonlinearities should be relatively small. We then examine, in our Monte Carlo study in Section 5, whether the option's gamma, which should proxy for the convexity of the option's payout, has any explanatory power for average returns. We then repeat this procedure in actual data. In both settings, we find the resulting bias to be miniscule.

3.4.2 Measurement errors

A bigger potential issue is the presence of measurement errors in option stock prices. As is standard, we proxy for "true" option prices by using the average of the bid and ask prices. Unfortunately, the difference between bid and ask prices of some option contracts can be very large, and if we assume that the "true" value of the option is merely within these quotes then errors in our midpoint estimate of the true price may be very large as well.

As Blume and Stambaugh (1982) point out, any error in the observed price will cause an upward bias in observed average returns. For instance, if f_t is the true option price and the observed price is equal to

$$\widehat{f_t} = f_t(1 + \delta_f), \tag{24}$$

then Blume and Stambaugh demonstrate that expected returns computed from observed prices are given by

$$\mathrm{E}\left[\frac{\widehat{f}_{t+1}}{\widehat{f}_t} - 1\right] \approx \mathrm{E}\left[\frac{f_{t+1}}{f_t} - 1\right] + \mathrm{Var}(\delta_f). \tag{25}$$

Matters are more complicated when one considers the bias in the return on a delta-hedged portfolio. In particular, if we are hedging an option contract by selling Δ shares financed at the risk-free rate, then the observed hedged return is equal to

$$\left(\frac{\widehat{f}_{t+1}}{\widehat{f}_t} - 1\right) - \frac{\Delta(\hat{S}_t, \sigma)\hat{S}_t}{\widehat{f}_t}\widehat{R}_{t+1} = \left(\frac{\widehat{f}_{t+1}}{\widehat{f}_t} - 1\right) - \beta_S(\hat{S}_t, \widehat{f}_t, \sigma)\widehat{R}_{t+1} \tag{26}$$

where $\hat{S}_t$ is the observed price of a share of the underlying stock, $\widehat{R}_{t+1}$ is the observed excess return of the stock and $\beta_S = \Delta(\hat{S}_t, \sigma)\hat{S}_t/\widehat{f}_t$ is the beta of the option with respect to the underlying stock.

Note that the observed hedge ratio $\beta_S(\hat{S}_t, \hat{f}_t, \sigma)$ is not generally equal to its true value. This is because any error in either the option or stock price will result in an error to the delta computed for the option and/or the prices used to normalize that delta.

In practice, we find that the bias in

$$\beta_S(\hat{S}_t, \widehat{f}_t)\widehat{R}_{t+1} \tag{27}$$

induced by errors in the option price is relatively small, and would be zero exactly in the case that stock returns had zero mean. Errors in the stock price have a much greater effect, in part because they bias the observed stock return, but also because they induce a nonzero covariance between the observed stock return and $\Delta(\hat{S}_t, \sigma)$. To understand where this covariance arises, consider the possibility that the observed stock price is above the true stock price ($\hat{S}_t > S_t$). This will cause a call option to appear more in-the-money than it actually is, which causes an upward bias in the option delta, which in turn implies that more shares of stock are short sold to delta-hedge when the observed stock price is above the actual stock price and, as a result, the delta-hedged return becomes positively biased.

To deal with this bias, we derive a close approximation for it along the same lines as Blume and Stambaugh (1982). In the discussion immediately below, we assume that stock price errors affect $\Delta(\hat{S}_t, \sigma)$ only directly through their effect on the moneyness of the option. More realistically, stock price errors have an additional effect that arises through the errors in implied volatilities that they induce, i.e. $\sigma = \sigma(\hat{S}_t)$. We derive the bias in this more realistic setting in Appendix E and present the result later in this section. Our derivations here are sufficient, however, to establish the intuition behind all of our more complex bias computations.

Assume that stock prices are observed with some error, so that $\hat{S}_t = S_t\,(1 + \delta_s)$, where δ_s has a mean of zero. Observation error may also be present in the stock price at time $t+1$, but as it has no effect on the bias, we will ignore it here.

Under these conditions, the bias is approximately equal to half of the second derivative of the return component, evaluated at $\hat{S}_t = S_t$, multiplied by the variance of the measurement error.[1] This is

$$\frac{1}{2}\frac{\partial^2}{\partial S_t^2}\left(\beta_S(S_t)\left(\frac{S_{t+1}}{S_t} - 1\right)\right)\mathrm{Var}(S_t\delta_s)$$

$$= \left(\frac{1}{2}\beta''_S(S_t)S_t^2\left(\frac{S_{t+1}}{S_t} - 1\right) - \beta'_S(S_t)S_{t+1} + \beta_S(S_t)\frac{S_{t+1}}{S_t}\right)\mathrm{Var}(\delta_s)$$

This is the bias *conditional* on S_{t+1}. The unconditional bias is

$$= \left(\frac{1}{2}\beta''_S(S_t)S_t^2\mu - \beta'_S(S_t)S_t(1+\mu) + \beta_S(S_t)(1+\mu)\right)\mathrm{Var}(\delta_s),$$

where μ is the mean return on the underlying stock.

Note that in most cases μ is small, which suggests that we can approximate the bias by setting μ to zero,

[1] Without loss of generality, we assume in this derivation that the stock dividend yield and the riskless rate are zero.

which results in

$$[-\beta_S'(S_t)S_t + \beta_S(S_t)]\text{Var}(\delta_s) \tag{28}$$

Note that the term $\beta_S(S_t)\text{Var}(\delta_s)$ is exactly the bias that would follow from Blume and Stambaugh (1982) in the case in which $\beta_S(S_t)$ were fixed. The term $-\beta_S'(S_t)S_t\text{Var}(\delta_s)$ therefore must be included to account for the nonzero covariance between $\beta_S(S_t)$ and the stock return that results from errors in the time-t stock price. This term is not small and cannot be ignored. Finally, to compute the bias we must find the first derivative of β_S, which for regular delta-hedged returns is[2]

$$\beta_S'(S_t) = \frac{\Gamma(S_t)S_t + \Delta(S_t)}{f_t} \tag{30}$$

The above result is valid only when implied volatility is known. If volatility is not observed and must be inverted from option and stock prices via the function

$$\sigma(\hat{S}_t, \hat{f}_t),$$

then an additional dependence on the stock price error is introduced, causing β_S' to take a different form. In particular, if we take the total derivative of

$$\beta_S(S_t) = \frac{\Delta(S_t, \sigma(S_t))S_t}{f_t}, \tag{31}$$

then we obtain

$$\beta_S'(S_t) = \frac{\Gamma(S_t, \sigma(S_t))S_t}{f_t} + \frac{\Delta(S_t, \sigma)}{f_t}\frac{d_1}{\sigma\sqrt{\tau}} \tag{32}$$

The form of the bias in (28) is the same as before.

This is the bias of a equally-weighted portfolio of delta-hedged options. Our results, however, are primarily based on hedging with the *total* delta, as defined in equation (5), which results in an option beta of

$$\beta_T(S_t) = \frac{\left(\Delta(S_t) + \nu(S_t, \sigma(S_t))\frac{\omega\rho}{\sigma(S_t)S_t}\right)S_t}{f_t} \tag{33}$$

$$= \beta_S(S_t) + \frac{\nu(S_t, \sigma(S_t))\omega\rho}{\sigma(S_t)f_t}$$

[2] Note that previous calculations treated the option price as a function of S, but this is not appropriate if we don't think that the option price is affected by stock price measurement error. The old result was:

$$\beta_S'(S_t) = \frac{\Gamma(S_t)S_t f(S_t) + \Delta(S_t)f(S_t) - \Delta(S_t)^2 S_t}{f(S_t)^2} \tag{29}$$

The form of the bias, as in (28), is unchanged, except that $\beta'_S(S_t)$ is replaced by $\beta'_T(S_t)$, which we show in the appendix is given by

$$\beta'_T(S_t) = \beta'_S(S_t) + \frac{\omega\rho}{\sigma(S_t)^2 f_t}\left(\frac{\nu(S_t,\sigma)\sigma(S_t)}{S_t}\left(1 - \frac{d_1}{\sigma\sqrt{\tau}}\right) + \Delta(S_t,\sigma(S_t))\left(1 - d_1 d_2\right)\right) \tag{34}$$

Rather than trying to adjust for this bias, Blume and Stambaugh's suggestion was to examine value-weighted instead of equally-weighted returns. While option prices can be observed readily, the zero net supply of each option contract makes the concept of market value ill defined. In some cases, we attempt to capture some of the benefits of Blume and Stambaugh's insight by weighting options according to the relative option price, or the price of the option divided by the price of the underlying stock. In this case our portfolio return is an equal weighted average of

$$w(\hat{S}_t)\left(\frac{\widehat{f}_{t+1}}{\widehat{f}_t} - 1\right) - w(\hat{S}_t)\beta(\hat{S})\left(\frac{S_{t+1}}{\hat{S}_t} - 1\right) \tag{35}$$

where

$$w(\hat{S}_t) \propto \frac{\hat{f}_t}{\hat{S}_t} \equiv \frac{k\hat{f}_t}{\hat{S}_t}$$

and the "weights" sum to N (not 1), the number of assets in the portfolio. We will assume that the constant of proportionality k is known exactly, which would approximately be the case in a large portfolio if measurement errors are independent.

The first component,

$$w(\hat{S}_t)\left(\frac{\widehat{f}_{t+1}}{\widehat{f}_t} - 1\right) = \frac{k\hat{f}_t}{\hat{S}_t}\left(\frac{\widehat{f}_{t+1}}{\widehat{f}_t} - 1\right) = \frac{k}{\hat{S}_t}\left(\widehat{f}_{t+1} - \widehat{f}_t\right)$$

is very close to unbiased (exact if the true f_t is a martingale). The second component in the difference,

$$w(\hat{S}_t)\beta(\hat{S}_t)\left(\frac{S_{t+1}}{\hat{S}_t} - 1\right),$$

has a bias that is approximately equal to

$$\left[-k\hat{f}_t\beta'(S_t) + 2\frac{k\hat{f}_t}{S_t}\beta(S_t)\right]\mathrm{Var}(\delta_s),$$

where $\beta(S_t)$ is defined as above (either as β_S or β_T). Thus, the Blume and Stambaugh suggestion of weighting by value succeeds at eliminating bias in the option return component, but bias adjustment is still necessary as a result of the delta hedge.

We do not observe $\text{Var}(\delta_s)$ or $\text{Var}(\delta_f)$, and therefore use proxies for both in the Fama-MacBeth regressions. We assume that $\text{Var}(\delta_f)$ is proportional to the square of the relative option spread (the spread divided by the midpoint), and therefore use the mean of the square of the relative options spreads in the portfolio as a dependent variable in the second-pass Fama-MacBeth regression. In the simulations, we assume that $\text{Var}(\delta_s)$ is proportional to the square of the effective stock spread. Consequently, we use the mean across all options in a portfolio of $[-\beta'_S(S_t)S_t + \beta_S(S_t)]$ multiplied by the squared stock spread as a dependent variable in the second-pass Fama-MacBeth regression. In the actual data, we assume that $\text{Var}(\delta_s)$ is given by a linear function of the squared effective spread. We therefore use both the mean of $[-\beta'_S(S_t)S_t + \beta_S(S_t)]$ and the mean of $[-\beta'_S(S_t)S_t + \beta_S(S_t)]$ multiplied by the squared stock spread as dependent variables in the Fama-MacBeth regressions performed on the actual data. We show by means of the Monte Carlo study described in Section 5 that adding the bias adjustments above properly cleans the effects of bid-ask spreads in the estimation of the volatility-risk premium.

3.4.3 Censoring

A third major potential concern in the implementation of the Fama-MacBeth procedure is related to the fact that implied volatilities cannot be computed for options whose prices violate arbitrage bounds. Without an implied volatility, it is impossible to compute the hedge ratios necessary for the computation of hedged returns. Consequently, these options cannot, without some additional assumptions, be used in our Fama-MacBeth procedures.

The usual procedure adopted in the literature is to discard options that violate arbitrage bounds.[3] However, our Monte Carlo simulation results below show that discarding options without implied volatility can result in serious biases. The reason is that arbitrage bounds tend to be violated when measurement error in prices is positive, but not when it is negative. By systematically excluding positive pricing errors when hedged option positions are put on, average measured returns will be biased upwards.

To get around this issue we include options even when they violate arbitrage bounds. We do so by computing the hedge ratios using the implied volatility from another option on the same firm. For a put (call) that is missing an implied volatility, our first choice is to use the call (put) with the same maturity and moneyness. If that value is unavailabe, then we use the implied volatility on the same contract from the previous day (or before then if necessary). This implied volatility is then used to calculate the delta and the vega of the problematic contract.

[3] See for instance Santa-Clara and Saretto (2005).

4 Data and summary statistics

The options data are from the Ivy database. The Ivy options database provides end-of-day bid and ask quotes, open interest, trading volume, implied volatility, and option hedge ratios for the US equity and index options market. We use options data from January 4, 1996, through December 31, 2005. The implied volatility and option hedge ratios are calculated with the Black and Scholes (1973) model for European options, or the Cox, Ross, and Rubinstein (1979) binomial tree for American options. All options on individual stocks are American. The dataset also includes a volatility-surface file that contains the interpolated volatility surface for each security on each day using a kernel-smoothing algorithm. The volatility surface provides the implied volatility for each standardized option with various times-to-maturity and deltas. The only variable used from the volatility surface file is the thirty-day implied volatility with delta equal to 0.5. This interpolated implied volatility is used to calculate the moneyness of options[4].

Stock prices and returns are from the CRSP database. Our sample includes all listed stocks with options over sample period. We merge our CRSP stock dataset and option dataset by matching securities' CUSIPs. Because our main analysis also considers the transaction costs in the stock market, we utilize the stock transaction database from TAQ. We measure the transaction cost of a stock in a given day as half of the daily effective spread. To estimate the daily effective spread we use Lee and Ready (1991) algorithm to match each transaction in the TAQ database to its corresponding quote, the effective bid-ask spread is then calculated as the ratio of the quoted bid-ask spread and the transaction price. Daily estimates are obtained as simple averages throughout the day. The procedure filters used in the TAQ database are the same as those described in Duarte and Young (2007).

We delete a series of observations to eliminate possible data errors. Appendix C describes the data cleaning procedures. The resulting dataset has 5,156 different stocks, more than 99 million daily call option observations and more than 82 million daily put option observations over the period between 1996 and 2005. Around 17 percent of these observations do not have implied volatility computed by Ivy's database because Ivy does not compute implied volatilities of options with vega smaller than 0.5 or with mid-point of the bid-ask price smaller than the option's intrinsic value. In addition, Ivy does not display implied volatilities when the algorithm to calculate implied volatilities fails. After using the procedure described in Section 3.4 to compute the implied volatilities and hedge ratios of these options, less than two percent of the options in the final database do not have computed implied volatility.

Table 1 presents the cross-sectional distribution of a series of variables of interest. The implied volatility (σ) in this table is the 30-day implied volatility of the options with delta equal to 50%. This implied volatility

[4] We define moneyness as $\ln(K/S)/(\sigma\sqrt{T})$ where K is the option exercise price, σ is the stock volatility and T is the time-to-maturity.

is between 20 and 104 percent for the majority of stocks. The distribution of estimates of the correlation between daily returns and changes in implied volatility (ρ) is in general negative with a median value of -26 percent. The distribution of estimates of the annualized volatility of volatility (ω) based on the daily standard deviation of the changes in implied volatility indicates that the median value of ω is around 84 percent. The cross-sectional distribution of the betas of the total-delta-hedged options with respect to the market and with respect to the market volatility is also presented. The beta of the total-delta-hedged options with respect to the market is in general negative with a median value of -0.06, the beta with respect to the volatility of the market is in general positive with a median value of 0.23.

The results in Table 1 also indicate that the total-delta-hedge method has essentially the same performance as does the standard delta method. The first row of the second panel in Table 1 displays the distribution of the standard deviation of the daily returns of delta-hedge method estimated for an entire year for each individual stock. These standard deviations are quite large with a mean value of 17%. These large standard deviations of daily delta-hedged returns are caused by the fact that the standard deviation of daily returns of out-of-the-money options is quite large, and because the bid-ask spreads of these options is so large in relation to their prices that a significant part of the price variation of these options is related to microstructure noise. The second row of the second panel in Table 1 displays the distribution of the difference between the standard deviation of the daily returns of the delta-hedge method and the standard deviation of the daily returns of the total-delta-hedged returns. These differences are in general negative with a median value of minus five basis points, indicating that the total-delta hedging method has slightly better performance than the traditional delta-hedge method.

The mean returns of calls and puts displayed in Table 2 indicate that the possible biases caused by censoring options without implied volatility and by large bid-ask spreads of options are large. Table 2 presents the mean and t-statistics of daily returns, absolute bid-ask spreads and relative bid-ask spreads of calls and puts across different times-to-maturity and moneyness. The mean returns of options in the first panel of Table 2 are based on the usual procedure of eliminating all options without implied volatility. The second panel of Table 2 presents the mean returns of calls and puts using all the options in the database including those without implied volatility. Note that the mean returns of calls without censoring are in general positive, indicating that the censoring of options without implied volatility can lead to the inaccurate conclusion that the expected returns of in-the-money calls are negative. The second panel of Table 2 also indicates that the mean return of the out-of-the-money puts is positive, however the large bid-ask spreads of out-of-the-money puts displayed in the third panel of Table 2 suggests that the positive mean returns of out-of-the-money puts may result from Blume and Stambaugh's biases described in Section 3.4. The Monte Carlo simulations below confirm that the signs of the mean returns of both calls and puts can be explained by censoring and Blume and Stambaugh's biases.

Recall that the market one-vega P&L defined in Section 3.2 is estimated by the mean return of a portfolio of options mimicking the return of a long position on an option with vega equal to one, financed at the risk-free rate. Figure 1 plots the time series of the market one-vega P&L along with the time series of level of the S&P 500 index. The market one-vega P&L is usually negative with a mean value in the period of -4.6 basis points. As previously mentioned, the mean value of the market one vega is also an estimate of the volatility risk premium. Therefore the time series of the market one vega implies that the volatility risk premium is negative and economically significant with a value of -4.6 basis points per day, or minus eleven percent per year.

Figure 1 also reveals that the market one-vega P&L has positive spikes as high as 15 percent on some days. Some of the positive spikes in the market one-vega P&L are consistent with periods when the market was under distress, such as during the Long Term Capital Management debacle in October 1998 and September 2001. Overall, we interpret the time series of the market one-vega P&L as being consistent with the idea that long positions in delta-hedged S&P 500 options provide an insurance against market distresses.

5 Monte Carlo analysis

We use a Monte Carlo Analysis to show that our design of the Fama-MacBeth regressions results in reasonable estimates of the volatility risk premium. We simulate 300 ten-year time series of daily prices of puts, calls and stocks. The simulated calls and puts have time-to-maturity between ten and 65 days and moneyness between minus three and three. These options are written on 1,000 different stocks. For each simulation path, we estimate a Fama-MacBeth procedure as described in Section 3.3. We also calculate average returns of portfolios of options with different moneyness and times-to-maturity. Because the simulations are computationally intensive, we distribute them across more than one hundred PCs.

The underlying model in our Monte Carlo simulations is a model with stochastic volatility[5]. The underlying model is a parameterization of the Heston's model and it is described in Section 5.1. In one set of Monte Carlo simulations, we assume that the price of volatility risk is equal to $\lambda_{2,t}^m = -0.1$. This negative price of volatility risk is consistent with the result in Coval and Shumway (2001), that selling delta-hedged options results in positive abnormal returns. In a second set of Monte Carlo simulations, we assume that volatility is not priced $\lambda_{2,t}^m = 0$. This second set of Monte Carlo simulations is interesting because in the Fama-MacBeth regression we test the hypothesis that $\lambda_{2,t}^m = 0$ and hence these second set of Monte Carlo

[5] A separate set of Monte Carlo simulations with the Black Scholes model is also executed. The results are essentially the same as those reported herein and therefore are not diplayed.

simulations are useful for analyzing how often our Fama-MacBeth regressions would find a volatility risk premium when there is none.

Recall that the relationships between delta-hedged option returns and prices of risks are based on continuous hedging, while we work with daily hedging rebalance in our empirical applications. To quantify the magnitude of the error caused by the fact that we do not continuously delta-hedge the options in our application, we run the Fama-MacBeth regressions in the Monte Carlo simulations with an additional variable on the right hand side: The "adjusted gamma" (the gamma of the option divided by its price multiplied by the squared stock price.). The rationale for this procedure is that if the daily delta-hedge procedure is a poor approximation for the continuous delta-hedging procedure then options with large adjusted gamma would have higher expected returns. (For proof see Appendix D.)

To understand the effects of measurement errors in our empirical analysis, we generate two sets of simulated prices: One set of simulations is based on prices without errors and a second set of simulations is based on prices with errors. In the simulations with errors, each simulated option price is assumed to equal the "true" value from the Heston model, plus some measurement error arising from nonzero bid-ask spreads. Errors are also introduced into stock prices by perturbing the simulated price processes with serially uncorrelated errors.

Because option bid-ask spreads are highly related to maturity and moneyness, we calibrate our simulation to match the averages and standard deviations of the bid-ask spreads as a fraction of the bid-ask midpoint that we observe for the different option categories in the data. Table 2 displays these values for a variety of maturity and moneyness buckets. The overall pattern that emerges is that OTM options have far greater spreads than ITM options, both for puts and for calls. In addition, both the level and variability of spreads show a moderate tendency to decline with option maturity.

To simulate option bid-ask spreads, each stock is assigned a constant measure of the liquidity, η_i, of all the options written on it. This value is drawn from a standard normal distribution and does not change over time. The maturity and moneyness of each option contract is then used to interpolate the values in the first panel of Table 3. The bid-ask spread of the contract is then computed as $e^{M+S\eta_i}$. Note that both M and S change through time with the maturity and moneyness of the contract, but η_i remains the same. This induces cross-sectional variation in option spreads, even for similar contracts, while at the same time making the spreads of a given contract somewhat persistent through time.

In our option price data, we observe bid-ask spreads but not transaction prices. We view the spread as providing bounds on the true value of the option, which is assumed, on average, to equal the bid-ask midpoint. Option price errors are therefore no greater than plus or minus half of the bid-ask spread. We capture these assumptions in the triangular error distribution[6], as illustrated in Figure 2.

Table 2 shows that the relative bid-ask spreads of options can be as large as one hundred percent. These large bid-ask spreads can cause biases in expected returns that cannot be properly captured by the approximations derived in Section 3.4.2 because those approximations were based on a relatively simple second-order expansion around a zero bid-ask spread. We therefore analyze a set of simulations in which we filter out options with bid-ask spreads above 25% of the mid price. To avoid any spurious correlations that this filtering may cause on our return calculations at time t, we base this filter on the spreads observed at time $t-2$. We hereafter call the procedure of deleting these options*filtering.*

Error in *hedged returns* is also induced by mismeasurement of stock prices. Since our stock price data consists solely of transaction prices, it is difficult to put bounds on these errors. We assume, therefore, that stock price errors are normally distributed with a mean of zero and a standard deviation that is proportional to the true price. The constant of proportionality is drawn at random for each stock from a uniform distribution between 0.001 and 0.005. These values are close to the effective stock bid-ask spreads in the data.

The effects of censoring options that violate non-arbitrage bounds is analyzed through two sets of Monte Carlo simulations: In one set, options violating arbitrage bound or with vega smaller than 0.25 are not considered in the empirical analysis[7]. We call the procedure of deleting these options *censoring*. In another set of simulations, the implied volatility of options is calculated by replacing the implied-volatility of options that violate non-arbitrage bounds as described in Section 3.4.

5.1 A parameterization of the Heston model with aggregate risk

The Heston model is a special case of the model discussed in the previous section. It is traditionally written using a process for stochastic variance rather than volatility. Under the risk-neutral measure, price and variance dynamics for a security i are

$$\begin{aligned} \frac{dS_t^i}{S_t^i} &= (r_t - q_t^i)dt + \sqrt{V_t^i}dB_{1,t}^{i*} \\ dV_t^i &= \kappa_V^i \left(\overline{V}^i - V_t^i\right) dt + \omega_V^i \sqrt{V_t^i} \left(\rho^i dB_{1,t}^{i*} + \sqrt{1-(\rho^i)^2} dB_{2,t}^{i*}\right) \end{aligned} \tag{36}$$

where $B_{1,t}^{i*}$ and $B_{2,t}^{i*}$ are Brownian motions under the risk-neutral measure.

We use the same model for both the market index and for each individual equity, where the subscript m denotes the market index. We assume that all covariances between the market index and the individual

[6]We also run simulations with uniformly distributed errors. The results are qualitatively similar to the ones with triangular distributed errors.

[7]Recall that Ivy database does not calculate implied volatility of options with vega smaller than 0.25.

securities are described by

$$dB_{j,t}^{i*} = \xi_j^i dB_{j,t}^{m*} + dZ_{j,t}^{i*} \tag{37}$$

for $j \in \{1, 2\}$, where all of the Z^* processes are independent zero-drift Brownian motions with instantaneous volatilities of $\sqrt{1 - (\xi_j^i)^2}$. That is, in the parameterization of the Heston model, we assume that $\xi_{12}^i = \xi_{21}^i = 0$ in Equation (9). These assumptions imply a single-factor model for equity returns. They also imply a single-factor structure in the part of each firm's volatility risk that is orthogonal to that firm's stock price. While this is not the most general covariance structure possible, it should capture most of the comovement in returns and volatilities at least approximately.

Let $B_{1,t}^m$ and $B_{2,t}^m$ represent Brownian motions under the physical probability measure. We assume that $B_{1,t}^m$ and $B_{2,t}^m$ are priced risks and that all of the Z processes are unpriced. Specifically, we assume that

$$dB_{j,t}^m = dB_{j,t}^{m*} + \lambda_j dt \tag{38}$$

for $j \in \{1, 2\}$. This specification of the prices of risk implies that instantaneous expected returns, under the objective measure, are given by

$$\mathrm{E}\left[\frac{dS_t^i}{S_t^i}\right] = \left(r_t - q_t^i + \lambda_1 \xi_1^i \sqrt{V_t^i}\right) dt. \tag{39}$$

The expected return of the market is obtained by setting $\xi_1^i = 1$. The specification in Equation (38) also implies that the drift of the variance process, under the objective measure, is

$$\kappa_V^i \times \overline{V}^i + [\lambda_1 \xi_1^i \rho^i + \lambda_2 \xi_2^i \sqrt{1 - (\rho^i)^2}] \times \omega_V^i \sqrt{V_t^i} - \kappa_V^i \times V_t^i \tag{40}$$

The admissibility of the V_t^i process under the true probability measure is discussed in Duarte (2004). To generate a positive equity premium, it must be the case that $\lambda_1 > 0$. The typically negative correlation between returns and volatility shocks requires that $\rho^i < 0$. The price of volatility risk can be either positive or negative; the results in Coval and Shumway (2001) and Bakshi and Kapadia (2003) suggest that $\lambda_2 < 0$.

We hedge options based on their Black and Scholes hedge ratios even though option prices are generated by the Heston model. The justification for this hedging procedure is twofold: First, in the Fama-MacBeth regression with actual data, we will use hedge ratios that are computed by the Ivy options database. Second, the use of Black and Scholes hedge ratios is much less computationally intensive than the use of Heston's hedge ratios.

The parameters used to simulate from the Heston model are described in Table 3. The parameters driving the risk-neutral process for the market variance are taken from Jones's (2003) estimates from the 1988-2000 sample period. The risk-premia parameter λ_1 produces an equity premium of about 4% per year, while the parameter λ_2 is either taken to equal zero or -0.1. The parameters displayed in Table 3 result in a volatility risk premium of -0.01[8]. This risk premium implies that the returns of delta-hedged at-the-money one-month straddles are around $-60bp$ per day, which is close to the value reported by Coval and Shumway (2001).

The parameters of the firm-level volatility processes are harder to estimate, and we know of no method that can produce accurate (e.g. maximum likelihood) estimates from firm-level data without incurring extreme computational costs. We therefore proceed more casually in selecting these parameters, starting with the assumption that all firm-level parameters are random i.i.d. draws from uniform distributions. While the specific distributions assumed are in Table 3, the parameters tend to give individual equities greater but slightly less persistent volatility with less negative skewness than the market volatility. In the absence (to our knowledge) of any empirical evidence on the cross-sectional distribution of the correlations between market and stock *variance* processes, we assume this distribution to be the same as the distribution of the correlations between market and stock *price* processes.

5.2 Simulation results

The simulations indicate that not controlling for measurement errors or for violations of arbitrage bounds can seriously mislead researchers performing empirical work with options. Table 4 displays the mean simulated returns of calls and puts when the price of volatility risk is negative. The first panel of this table displays the mean returns of calls and puts when the simulated bid-ask spread of options is zero. Note that in-the-money call options have positive mean returns and out-of-the-money call options have negative mean returns[9]. Put options, on the other hand, have negative mean returns. The second panel of Table 4 displays the mean simulated returns of calls and puts when bid-ask spreads are different from zero and options violating arbitrage bounds are not thrown away. Note the dramatic difference between the mean returns displayed in the first and second panel of this table. When bid-ask spreads are not zero, the mean returns of calls are always positive, and in fact the simulated mean returns of deep-out-of-the-money calls are quite large. Deep-out-of-the-money puts, on the other hand, have positive and large mean returns when measurement error is introduced in the simulation. The positive and large returns of calls and puts are consistent with a very large Blume and Stambaugh's bias in the mean returns of options. The third panel of Table 4 displays

[8] To see this note that the volatility risk premium in Equation (11) is equal to $\omega^m \sqrt{1-(\rho^m)^2}\lambda_2^m$ and by Ito's lemma $\omega_t^m = 0.5 \times \omega_V^m$.

[9] This result is consistent with those in Ni (2006) and are a consequence of a negative price of volatility risk.

the mean returns of calls and puts when the options have a non-zero bid spread and censoring is applied. Note that censoring in general decreases the mean returns of options. Again, the difference between the results in the third panel and the results in the second panel is dramatic, indicating that the effects of censoring options without implied volatility in the Ivy database are economically significant.

It is also interesting to note the similarity between the mean option return of the actual option sample and the mean returns of the simulated sample. The third panel of Table 4 has a similar pattern of mean returns as those displayed in the first panel of Table 2. Specifically, the results in these panels indicate that censoring causes the mean returns of deep-in-the-money call options to be negative in both the actual and simulated data. Also, the second panel of Table 2 is similar to the second panel of Table 4, which indicates that the positive returns of deep-out-of-the-money options is driven by their large relative bid-ask spreads in both the actual and simulated samples.

The results of the Fama-MacBeth regressions on the simulations without measurement errors indicate that the discretization bias is small. The first panel of Table 5 displays the results of Fama-MacBeth regressions in the simulations in which both stocks and options are observed without error. The results of the first Fama-MacBeth regression in this panel indicate that our estimated volatility risk premium is -0.005. This is closer to zero than the actual simulated volatility risk premium which is -0.01. The bias in this estimation is not surprising, granting the classic errors in variables problem, which causes a shrinking bias in the estimated risk premium. The third Fama-MacBeth regression in this panel also estimates the coefficient on the option gamma (γ). This is done to analyze the effect of discretization in our estimation, because, as previously mentioned, we would expect to see a statistically significant coefficient on gamma (γ) if the discretization were causing a large bias in the estimation. Note however that the estimated coefficient on gamma is not statistically significant in these Monte Carlo simulations. Indeed the average t-statistic of this coefficient is only -0.21.

The simulation of the Fama-MacBeth regressions indicate that we can properly estimate the volatility risk premium providing that we properly control for measurement errors in the options and stock market. The second panel of Table 5 displays the results of Fama-MacBeth regressions on simulations with measurement errors and no filtering of options. The first Fama-MacBeth regression in this table clearly shows the effect of measurement error on the estimation of the volatility risk premium, because the estimated volatility premium in this regression is not only large but also has the wrong sign. Indeed, the estimate volatility risk premium without controls for measurement errors is 3.6%. Controlling for options and stock price measurement errors change this estimation substantially. For instance, in the fourth Fama-MacBeth regression described in Table 5, the estimated volatility risk premium is -0.009, and recall that the simulated value is -0.01. The third panel of Table 5 displays the results of the Fama-MacBeth regressions in the simulations, where there are measurement errors and filtering. The first regression in this panel indicates that filtering has a substantial effect in reducing the effects of measurement errors in the estimation of risk premium. To see this compare the volatility risk premium estimated in the second regression of the second panel (0.39) with the volatility risk premium estimated in the first regression of the third panel (-0.003). Finally, the last regression in the third panel of Table 5 indicates that filtering, along with introducing controls for option and stock bid-ask spreads essentially cleans all of the biases related to measurement errors. In fact, the estimated risk premium in this regression is the same as the one estimated in the case where there no measurement errors in stocks and options (see the first panel of Table 5).

The Monte Carlo simulations also indicate that our Fama-MacBeth procedure can be effectively used to test for the presence of a volatility risk premium if measurement errors are controlled through additional right-hand side variables in the Fama-MacBeth regressions, and options with large bid-ask spreads are filtered. Table 6 displays the results of Fama-MacBeth regressions when the volatility risk premium is zero. The first regression in this table does not control for measurement errors and as a result, the estimated volatility risk premium is positive and large. The second regression in Table 6 controls for measurement errors through the inclusion of bid-ask spread related variables. In this case, the estimated volatility risk premium is on average -0.004 with an average t-statistic of -2.42. These results therefore indicate that controlling measurement errors only with bid-ask spread variables in the Fama-MacBeth regression would result in a frequent rejection of the hypothesis that the volatility-risk premium is zero when the volatility risk premium is in fact zero. (The probability of Type II error would be high.) The third regression in Table 6 controls for bid-ask spread biases through the inclusion of right-hand side variables and through filtering of options with large bid-ask spreads. In this regression, the average estimated volatility risk premium is zero with an average t-statistic of 0.06, indicating that filtering plays an important role in testing for the presence of a volatility risk premium.

6 Empirical results

6.1 An application to S&P 500 returns

To illustrate the effects of the measurement errors in the actual data, we calculate the mean returns of S&P 500 options with various times-to-maturity and moneyness. These mean returns are displayed in the second panel of Table 7. Note that the mean returns of call options on the S&P 500 are increasing with moneyness, in fact the deep-out-of-the-money short-term calls have a mean return close to 830 basis points per day. This result is consistent with Coval and Shumway (2001). To correct for the biases caused by measurement errors we run a Fama-MacBeth regression of S&P 500 option returns on the option spreads and use this regression to estimate the biases caused by measurement errors. These estimated biases are displayed in the third panel of Table 7 and they can be as high as 1413.8 basis points per day. The fourth panel of Table 7 displays the bias-adjusted mean returns of S&P 500. Note that this panel implies that the mean returns of call options are in fact decreasing with moneyness, which is exactly what we found in our simulations (see first panel of Table 4).

6.2 Constant risk premia

The first set of results on Fama-MacBeth regressions using actual data are displayed in Table 8. All the regressions in this and subsequent tables do not use censoring but do filter options with bid-ask spreads above 25%. For this table, portfolio returns are equal-weighted, and all results are based on OLS regression.

In addition to the two-factor specification discussed above, we also consider models that only include one of the two factors. The first three columns in the table correspond to specifications that exclude variables that control for measurement error-induced biases. Several findings here are notable. First, the estimated volatility risk premium is positive and significant for specification 1, which does not include a market return factor. When a market return beta is included, either in isolation or together with the volatility factor, it appears with a negative sign and is statistically significant. These results either contradict basic intuition or most of the previous literature examining equity index options.

Including bias control variables in specifications 4-6 results in a reversal of these findings. The price of volatility risk becomes negative, though only with marginal significance, and the price of market risk becomes insignificantly positive. The option spread control variable, which measures the average of the squared spread across all options in each portfolio, is highly significant, suggesting that the Blume and Stambaugh (1982) bias is an important component of *measured* average option returns. Biases related to measurement error in stock prices do not appear to be important here.

Two goodness-of-fit measures are included in the table. The first is the cross-sectional "R-squared", defined as one minus the ratio of the variance of the model-implied average returns to the variance of the actual average returns. The second (RMSE) is the square root of the average squared pricing error. Both measures are clearly improved by the addition of a volatility beta, suggesting the importance of the volatility factor in explaining realized returns. Somewhat unintuitively, the R-squared measure can go down (or the RMSE go up) when additional variables are added to the Fama-MacBeth regression. Because we are outside the simple linear regression framework, the expected result is not guaranteed. In some cases the cause is the presence of outliers in the bias control variables.

To check the robustness of these results, we re-estimate the last specification in Table 8 with three variations in methodology. First, we use relative price-weighted, rather than equal-weighted, portfolios. Next, we use weighted least squares (WLS) rather than OLS in the Fama-MacBeth cross-sectional regressions. Last, we combine both of these modifications. The results are presented in Table 9.

WLS regressions are appropriate in this case, at least as a robustness check, because of the substantial cross-sectional heteroskedasticity in the portfolios we consider. Deep out-of-the-money portfolios, for instance, are much more highly volatile than in-the-money portfolios. We therefore include WLS results by weighting each observation in a cross-sectional regression by the inverse of its factor model error variance.

Unfortunately, Table 9 contains few clear patterns to report. The negative volatility risk premium becomes more negative, and statistically significant, when using relative price-weighted portfolios, but the sign reverses for both specifications estimated using WLS. Furthermore, the two goodness-of-fit statistics (which are not themselves weighted) are much worse when prices of risk are computed using WLS regressions.

Figure 3 provides one possible explanation for these contradictory results. The figure is a scatter plot that shows the relationship between estimated volatility betas and average returns, both computed using the relative price-weighted portfolios. Points denoted "C" correspond to puts, while "P" refers to calls. Grey letters indicate portfolios of in-the-money or at-the-money options, while black letters denote out-of-the-money portfolios. The general pattern that emerges is that out-of-the-money options have low returns and high volatility betas. At the same time, there is a group of in-the-money options whose betas appear postively related to average returns.

In the OLS results, the variation in the observations in the bottom right part of the figure is dominant, and a highly significant negative relation is obtained. In the WLS results, however, those observations are downweighted due to their extreme volatility. Instead, the in-the-money options on the left side of the figure, with very low volatilities, are much more heavily weighted. Since those options exhibit a positive relation between volatility beta and average returns, the sign of the volatility risk premium becomes positive.

The last notable result from Table 9 concerns the controls for measurement error bias. We remarked above that the Blume and Stambaugh bias would likely disappear were we to examing relative price-weighted *unhedged* option returns. The second specification in the table shows that this result holds in practice as well, as the significance of the squared option spread variable vanishes. In its place, however, is an increase in the statistical significance of the biases that result from mismeasurement of the hedge ratio. Thus, bias controls are necessary in all cases, though the controls that are important depend on how the returns are weighted.

6.3 Time-varying risk premia

Though the sign of the unconditional volatility risk premium is difficult to determine, it is possible that nonzero conditional means are more easily detected. Figure 4 provides some visual evidence that this might be the case. The figure plots the time series of realized volatility risk premia ($\hat{\lambda}_{\mathrm{Vol}}(t)$) estimated from the Fama-Macbeth cross-sectional regressions run on equal-weighted portfolios using OLS. Serial correlation is clearly evident from the picture, as are a number of spikes corresponding to periods of major market turmoil, such as the market crash of October 1997 or the attacks of 9/11/2001.

A more formal analysis of serial correlation is provided in Table 10. In the column of results in the top panel of the table, $\hat{\lambda}_{\mathrm{Vol}}(t)$ is regressed on its own first lag. A significant positive autocorrelation estimate is obtained, with a moderate R-squared of .027. The next three columns address whether the same predictor can forecast realized market risk premia ($\hat{\lambda}_{\mathrm{M}}(t)$) or realized factor returns. The evidence is weaker for these variables.

Because a single lag of $\hat{\lambda}_{\mathrm{Vol}}(t)$ is undoubtedly noisy, the second panel of Table 10 instead uses a 22-day moving average of the predictor. The left-hand-side variables are unchanged. The result is that volatility risk premia seem to have some power to predict future realized market risk premia ($\hat{\lambda}_{\mathrm{M}}(t)$) and excess market returns, though evidence of the latter is not strong.

In this section we re-analyze the Fama-MacBeth regressions of the previous section too see if there are any discernable patterns in ex ante risk premia. Rather than looking at unconditional means of the $\hat{\lambda}_{\mathrm{Vol}}(t)$, we instead consider regressions of the form

$$\hat{\lambda}_{\mathrm{Vol}}(t) = a_{\mathrm{Vol}} + b_{\mathrm{Vol}} Z(t-1) + e(t), \tag{41}$$

where $Z(t)$ is some conditioning variable that is plausibly related to volatility (or market) risk premia. The realized risk premia $\hat{\lambda}_{\mathrm{Vol}}(t)$ are the same cross-sectional regression estimates used in the previous subsection.

Though theory does not provide a strong guide as to what conditioning variable to choose, most parametric models of option prices (e.g. Heston (1993)) assume that the asset's own volatility drives both the price of volatility and market risk. In the absence of obvious alternatives, we follow this literature and use the VIX index as a proxy for overall market volatility. The VIX is a measure of implied volatility for the S&P 500 index corresponding to a return horizon of one month. It is a "model-free" implied volatility of the type derived by Britten-Jones and Neuberger (2000), which relies on an many options with different strike prices but not any parametric model.

Table 11 contains results of the regression in (41) with VIX as the conditioning variable Z. Results assuming constant risk premia are repeated from previous tables. In addition to $\hat{\lambda}_{\mathrm{Vol}}(t)$, we also examine predictable variation in the market risk premia ($\hat{\lambda}_{\mathrm{M}}(t)$) and the intercepts ($\hat{\lambda}_0(t)$) of the Fama-MacBeth cross-sectional regressions.

The notation in the table follows from the fact that expected excess returns can now be represented (ignoring bias terms) as

$$(a_0 + b_0 \mathrm{VIX}(t-1)) + \beta_{\mathrm{Vol}} \left(a_{\mathrm{Vol}} + b_{\mathrm{Vol}} \mathrm{VIX}(t-1)\right) + \beta_{\mathrm{M}} \left(a_{\mathrm{M}} + b_{\mathrm{M}} \mathrm{VIX}(t-1)\right),$$

Thus, the rows labeled $\beta_{\mathrm{Vol}} \times \mathrm{VIX}(t-1)$ correspond to the coefficient b_{Vol}.

The first specification in Table 11 repeats a result from Table 8 showing an estimated positive price of volatility risk and a negative price of market risk when bias controls are excluded from the regression. In specification 2, in which volatility risk premia are allowed to vary with the level of the VIX, we see a significantly negative *constant* price of risk (-1.652×10^{-3}) and a significantly positive *variable* price of risk (0.849×10^{-2}). The same pattern obtains for specification 4, which includes our bias control variables. In this specification, the market risk premium also appears to be positively related with the VIX.

As before, robustness is a concern, so we repeat this last specification using a number of different alterations to our basic procedure. In Table 12, we run regressions using relative price-weighted portfolios and weighted least squares regressions. We also consider portfolios that are sorted by implied volatility or the ratio of implied to past realized 1-month volatility, both in addition to option maturity and moneyness. Finally, we repeat one of the specifications using regular delta-hedged, rather than total delta-hedged, returns. Our results are very robust to all of these choices. The volatility risk premium is increasing in the level of the VIX, and for all but one specification the market risk premium displays the same behavior.

6.4 Decomposing volatility betas

Our three-way portfolio sort picks up several different sources of variation in volatility betas. Sorting by maturity and moneyness separates out options whose contract characteristics, such as vega, make them more sensitive to volatility movements in general. These options should therefore be expected to vary more in response to movements in market volatility.

In many studies, usually of S&P 500 index options, these contract-level differences are the sole source of variation in betas. When these studies find that volatility risk is priced, a natural concern is that volatility risk sensitivity is merely proxying for an option characteristic that is related to expected returns for a completely different reason.

In our framework, variation in volatility betas also arises from differences in the covariances of the underlying stock with the market return and volatility factors. Hence, we can identify variation in option betas that is unrelated to contract specifications like maturity and strike price.

If volatility risk is the true driver of expected option returns, at least conditionally, then each source of variation in volatility betas should be equally important. If volatility betas are instead proxies for option contract characteristics that are related to expected returns for some other reason, then only variation due to contract type should appear to be priced, and variation due to differences across underlying assets should appear unimportant.

We propose an easy way to address this question. After estimating market and volatility betas in time series regressions, we average the volatility betas across all portfolios with the same maturity and moneyness (keeping puts and calls separate). Thus, all variation that is due to the third sort variable, which measures the covariance of the underlying stock's volatility with the market, is eliminated. These average volatility betas represent the "contract component" of the original betas. The "firm component" is then defined as the difference between the original portfolio beta and the contract component just defined. It captures the variation that is due to the nature of the underlying stocks.

Table 13 displays the results of Fama-MacBeth regressions in which volaility betas are decomposed into these two components. In addition, some results using regular or "total" betas are repeated from Tables 8 and 11. Specifications 1 and 2 examine risk premia that are assumed to be constant, while 3 and 4 allow risk premia to vary with the VIX index.

Examining the results in the second column, we see that the contract component of the volatility risk premium is significantly negative, while the firm component is positive with borderline significance. Since each of these terms should be identical to the "total" price of volatility risk in the first column, this result contradicts a fundamental implication of our theory of options pricing.

The third and fourth columns, where we are allowing volatility and market risk premia to vary with the VIX, present a very different picture. When the volatility beta is not decomposed, as we saw in Tables 11 and 12, it has a constant component that is significantly negative and a component related to the VIX that is significantly positive. The final column of Table 13 shows that both results hold for the "contract" and "firm" components of the volatility beta separately. Furthermore, the magnitudes of the "contract" and "firm" effects are roughly similar. These results suggest that volatility risk is not merely a proxy for contract characteristics.

7 Conclusion

This paper uses a large cross section of options on many different stocks to estimate the volatility risk premium. Though we find no reliable evidence that the price of volatility risk is nonzero on average, we document strong evidence of a conditional risk premium that varies positively with the overall level of market volatility. This result is robust to a variety of alternative portfolio construction methods, the use of WLS versus OLS, and whether or not bias control variables are included in the regression.

We analyze the effects of measurement errors, violations of arbitrage bounds, and discretization on the empirical work with options by means of a large Monte Carlo simulation exercise. Overall, our results indicate that discretization effects are small if daily data are used. On the other hand, the bias induced by censoring options that do not satisfy arbitrage bounds can be large, possibly resulting in biases in expected returns as large as several percentage points *per day* and possibly seriously misleading researchers working with options data. This bias arises when measurement errors cause the researcher to discard option prices that are too high, but not prices that are too low. We propose a simple procedure to retain these data points that has a large effect in simulations and in analysis of actual data.

Controlling for the bias in average returns due to measurement errors (bid-ask spreads) is perhaps even more important and of a magnitude sufficient to reverse some inferences about average option returns made in other studies. We show, for instance, that measurement error bias might be large enough to make out-of-the-money call options on the S&P 500 index appear to have positive and large average returns, when in fact the true returns may well be negative. Several different strategies that control for this bias are proposed by extending the work of Blume and Stambaugh (1982).

Finally, our simulations show that the Fama-MacBeth approach, when applied to option returns, can provide powerful and reliable inferences about the risk premia that drive the options market. A major advantage of this framework, besides its simplicity, is that it imposes minimal parametric structure, unlike most alternative approaches that rely on a complete specification of a stochastic volatility model for each underlying stock. In particular, unlike other parametric methods, our approach does not automatically generate a nonzero price of risk if average option returns are themselves nonzero – such an inference will only be drawn given evidence of a relation between volatility betas and average returns.

Appendix

A - Proof of Equation (3)

Ito's lemma implies:

$$\begin{aligned}\frac{df}{f} &= \mu_f dt + \frac{1}{f}(\frac{\partial f}{\partial S}S\sigma + \frac{\partial f}{\partial \sigma}\omega\rho)dB^i_{1,t} + \frac{1}{f}\frac{\partial f}{\partial \sigma}\omega\sqrt{1-\rho^2}dB^i_{2,t} \\ \mu_f &= \frac{1}{f}(\frac{\partial f}{\partial t} + \mu S\frac{\partial f}{\partial S} + \frac{1}{2}\sigma^2 S^2\frac{\partial^2 f}{\partial S^2} + \theta\frac{\partial f}{\partial \sigma} + \frac{1}{2}\omega^2\frac{\partial^2 f}{\partial \sigma^2} + \frac{\partial^2 f}{\partial \sigma \partial S}\omega\rho S\sigma)\end{aligned} \tag{42}$$

where the subscripts and superscripts on S, μ, σ, ρ and ω were removed for simplification.

Now, let's compute the excess returns of a derivative: Non-arbitrage implies that there is an equivalent martingale measure under which the dynamics of the stock price and of the volatility are:

$$\begin{aligned}\frac{dS}{S} &= (r-q)dt + \sigma dB^{*i}_{1,t} \\ d\sigma &= \upsilon dt + \omega\rho dB^{*i}_{1,t} + \omega\sqrt{1-\rho^2}dB^{*i}_{2,t}\end{aligned} \tag{43}$$

where $B^{*i}_{1,t}$ and $B^{*i}_{2,t}$ are Brownian motions. Non-arbitrage also implies that the price of any derivative satisfies the PDE:

$$\frac{\partial f}{\partial t} + (r-q)S\frac{\partial f}{\partial S} + \frac{1}{2}\sigma^2 S^2\frac{\partial^2 f}{\partial S^2} + \upsilon\frac{\partial f}{\partial \sigma} + \frac{1}{2}\omega^2\frac{\partial^2 f}{\partial \sigma^2} + \frac{\partial^2 f}{\partial \sigma \partial S}\omega\rho S\sigma = rf \tag{44}$$

Substituting the PDE above in the actual dynamics for f, we get:

$$\begin{aligned}\frac{df}{f} &= \frac{1}{f}\{(\mu - r + q)S\frac{\partial f}{\partial S} + (\theta - \upsilon)\frac{\partial f}{\partial \sigma} + rf\}dt + \\ &\quad + \frac{1}{f}(\frac{\partial f}{\partial S}S\sigma + \frac{\partial f}{\partial \sigma}\omega\rho)dB^i_{1,t} + \frac{1}{f}\frac{\partial f}{\partial \sigma}\omega\sqrt{1-\rho^2}dB^i_{2,t}\end{aligned} \tag{45}$$

Define the following prices of risk:

$$\begin{aligned}dB^*_{1,t} - dB_{1,t} &= \lambda^i_{1,t}dt = \frac{(\mu - r + q)}{\sigma}dt \\ dB^*_{2,t} - dB_{2,t} &= \lambda^i_{2,t}dt \\ \theta - \upsilon &= \omega\rho\lambda^i_{1,t} + \omega\sqrt{1-\rho^2}\lambda^i_{2,t}\end{aligned} \tag{46}$$

Substituting in Equations (45), we get:

$$\frac{df}{f} = \{\lambda^i_{1,t}\sigma\frac{S}{f}\frac{\partial f}{\partial S} + \omega\rho\lambda^i_{1,t}\frac{1}{f}\frac{\partial f}{\partial \sigma} + \omega\sqrt{1-\rho^2}\lambda^i_{2,t}\frac{1}{f}\frac{\partial f}{\partial \sigma} + r\}dt$$
$$+\frac{1}{f}(\frac{\partial f}{\partial S}S\sigma + \frac{\partial f}{\partial \sigma}\omega\rho)dB_{1,t} + \frac{1}{f}\frac{\partial f}{\partial \sigma}\omega\sqrt{1-\rho^2}dB_{2,t} \quad (47)$$

which implies that the expected excess return of the derivative is:

$$E[\frac{df}{f}] - rdt = [\{\sigma\frac{S}{f}\frac{\partial f}{\partial S} + \omega\rho\frac{1}{f}\frac{\partial f}{\partial \sigma}\}\lambda^i_{1,t} + \omega\sqrt{1-\rho^2}\frac{1}{f}\frac{\partial f}{\partial \sigma}\lambda^i_{2,t}]dt$$

B - Total-delta hedged portfolios

To hedge an option against stock price movements, one has to take a position on n stocks in a way such that stock price movements do not have an effect on the price of the hedged portfolio. That is, we choose n that isolates the value of the portfolio from movements in the Brownian motion $B_{1,t}$. Equation (47) makes clear that n is given by:

$$n = -(\frac{\partial f}{\partial S} + \frac{\partial f}{\partial \sigma}\frac{\omega\rho}{\sigma S}) \quad (48)$$

Consequently, the excess instantaneous return of a delta-hedged option when the position on the stock is financed at the risk-free rate is:

$$\frac{dH}{H} - rdt = \frac{df + ndS - n(r-q)Sdt}{f} - rdt = \frac{df}{f} + n\frac{S}{f}\frac{dS}{S} - n\frac{S}{f}(r-q) - rdt \quad (49)$$

where H is the value of the delta-hedged option portfolio. Substituting Equations (48) and (47) into Equation (49), we get:

$$\frac{dH}{H} - rdt = \omega\sqrt{1-\rho^2}\lambda^i_{2,t}\frac{1}{f}\frac{\partial f}{\partial \sigma}dt + \frac{1}{f}\frac{\partial f}{\partial \sigma}\omega\sqrt{1-\rho^2}dB_{2,t} \quad (50)$$

Ito's lemma also implies that the covariance between derivative returns and stock returns is:

$$cov(\frac{dS}{S}, \frac{df}{f}) = \frac{1}{f}\{\sigma^2 S\frac{\partial f}{\partial S} + \sigma\omega(\sigma)\rho\frac{\partial f}{\partial \sigma}\}dt \quad (51)$$
$$\beta_S = \frac{S}{f}\frac{\partial f}{\partial S} + \omega(\sigma)\rho\frac{1}{\sigma f}\frac{\partial f}{\partial \sigma}$$

where β_S is the beta of the derivative with respect to the stock.

C - Options database cleaning procedures

We do not use any options with non-standard settlement rules. All the options that have bid-prices equal to 998 or offer-prices equal to 999 are eliminated because Ivy used these as missing codes in some years. All the observations with negative bid-ask spreads or bid-ask spreads greater than five are eliminated. All calls or put options with delta smaller than minus one and greater than one are dropped from the database. All options with negative implied volatility are also dropped. (Ivy database uses -99.99 as a missing code in some observations.) At date t, we also eliminate all options with zero open interest at date $t-1$. This procedure helps to eliminate data errors that may be related to the lack of trading activity on some options. Options with a mid point of the bid and the ask prices below 50 percent of the intrinsic value or above 100 dollars of the intrinsic value are also eliminated. We eliminate all duplicated observations of the same type options with the same strike price and time to maturity on a given day. This guarantees that there is only one call or put with a given strike and maturity on any given day.

7.1 D - Daily delta-hedged returns

Let $H_{\tau,\tau+h}$ represent the value at time $\tau+h$ of a portfolio that was hedged at time τ. The portfolio is composed by a long position on a derivative ($f_{\tau+h}$), and n shares of the stock financed at the risk free rate (r). Assume that the underlying stock does not pay dividends. In this case $H_{\tau,\tau+h}$ is:

$$H_{\tau,\tau+h} \equiv f_{\tau+h} + nS_{\tau+h} - e^{rh}(f_\tau + nS_\tau) \tag{52}$$

Assume that the Black and Scholes model is used for the purpose of computing hedge ratios. In this case, $n = -\partial f/\partial S = -\Delta_\tau$. Substituting a Taylor approximation of $f_{\tau+h}\ e^{rh}$ in the expression of $H_{\tau,\tau+h}$, we get:

$$\begin{aligned} H_{\tau,\tau+h} \quad \approx \quad & f_\tau + \Theta_\tau h + \Delta_\tau \left(S_{\tau+h} - S_\tau\right) + \frac{1}{2}\Gamma_\tau \left(S_{\tau+h} - S_\tau\right)^2 \\ & -S_{\tau+h}\Delta_\tau - \left(f_\tau - S_\tau \Delta_\tau\right) - rh\left(f_\tau - S_\tau\Delta_\tau\right) \end{aligned} \tag{53}$$

$$H_{\tau,\tau+h} \approx \Theta_\tau h + \frac{1}{2}\Gamma_\tau S_\tau^2 \left(\frac{S_{\tau+h}}{S_\tau} - 1\right)^2 - rh\left(C_\tau - S_\tau\Delta_\tau\right) \tag{54}$$

Where $\Theta_\tau = \partial f/\partial t$ and $\Gamma_\tau = \partial^2 f/\partial S^2$. The Black and Scholes PDE is the following relation between Θ_τ, Γ_τ

and Δ_τ when the underlying security does not pay dividends:

$$\Theta_\tau + rS_\tau\Delta_\tau + \frac{1}{2}\sigma^2 S_\tau^2 \Gamma_\tau = rC_\tau \tag{55}$$

Substituting the PDE in the approximation for $H_{\tau,\tau+h}$ and taking expectations we have:

$$E[H_{\tau,\tau+h}] \approx \frac{1}{2}\Gamma_\tau S_\tau^2 \{E[\left(\frac{S_{\tau+h}}{S_\tau} - 1\right)^2] - \sigma^2 h\} \tag{56}$$

Using $\frac{S_{\tau+h}}{S_\tau} - 1 \approx \mu h$

$$E[H_{\tau,\tau+h}] \approx \frac{1}{2}\Gamma_\tau S_\tau^2 \mu^2 h^2 \tag{57}$$

Thus, the expected excess return on a delta-hedged option over period h is approximately

$$\frac{h^2}{2}\mu^2 \frac{S_\tau^2 \Gamma_\tau}{C_\tau} \tag{58}$$

Note that the expected return above depends on μ, which is difficult to proxy for. A simple strategy is to assume that μ is unknown but constant cross-sectionally. This implies that the discretization bias in delta-hedged returns is approximately proportional to

$$\frac{S_\tau^2 \Gamma_\tau}{C_\tau} \tag{59}$$

The bias in a portfolio is clearly no more than the weighted average of this bias. This holds when we are considering a portfolio of delta-hedged positions across different stocks, strike, or maturities.

E - Derivations of measurement error biases

Accounting for unknown implied volatility

The above result is valid only when implied volatility is known since $\Delta(S_t, \sigma(S_t))$ then has an additional dependence on the stock price. When implied volatility is unknown, β'_S takes a different form. In particular, if

$$\beta_S(S_t) = \frac{\Delta(S_t, \sigma(S_t))S_t}{f_t} \tag{60}$$

then

$$\beta'_S(S_t) = \frac{d\Delta(S_t, \sigma(S_t))}{dS_t}\frac{S_t}{f_t} + \frac{\Delta(S_t, \sigma(S_t))}{f_t} \tag{61}$$

$$= \left(\frac{\partial\Delta(S_t, \sigma(S_t))}{\partial S_t} + \frac{\partial\Delta(S_t, \sigma(S_t))}{\partial\sigma}\frac{d\sigma(S_t)}{dS_t}\right)\frac{S_t}{f_t} + \frac{\Delta(S_t, \sigma(S_t))}{f_t}$$

Note that $\partial\Delta/\partial\sigma$ is the *vanna* of the option, which is equal to

$$\frac{\nu(S_t,\sigma)}{S_t}\left(1-\frac{d_1}{\sigma\sqrt{\tau}}\right).$$

To find $d\sigma/dS_t$, note that σ solves $\hat{f}=f(S_t,\sigma)$. By the implicit function theorem,

$$\frac{d\sigma}{dS_t}=-\frac{\frac{\partial f}{\partial S_t}}{\frac{\partial f}{\partial\sigma}}=-\frac{\Delta(S_t,\sigma)}{\nu(S_t,\sigma)}$$

Putting everything together, we have

$$\begin{aligned}\beta_S'(S_t)&=\left(\Gamma(S_t,\sigma(S_t))-\frac{\nu(S_t,\sigma)}{S_t}\left(1-\frac{d_1}{\sigma\sqrt{\tau}}\right)\frac{\Delta(S_t,\sigma)}{\nu(S_t,\sigma)}\right)\frac{S_t}{f_t}+\frac{\Delta(S_t,\sigma(S_t))}{f_t}\\&=\frac{\Gamma(S_t,\sigma(S_t))S_t}{f_t}-\frac{\Delta(S_t,\sigma)}{f_t}\left(1-\frac{d_1}{\sigma\sqrt{\tau}}\right)+\frac{\Delta(S_t,\sigma(S_t))}{f_t}\\&=\frac{\Gamma(S_t,\sigma(S_t))S_t}{f_t}+\frac{\Delta(S_t,\sigma)}{f_t}\frac{d_1}{\sigma\sqrt{\tau}}\end{aligned}\tag{62}$$

Biases in total delta-hedged returns

Now consider the case in which the total delta is used for hedging, so that

$$\beta_T(S_t)=\frac{\left(\Delta(S_t)+\nu(S_t,\sigma(S_t))\frac{\omega\rho}{\sigma(S_t)S_t}\right)S_t}{f_t}\tag{63}$$

$$=\beta_S(S_t)+\frac{\nu(S_t,\sigma(S_t))\omega\rho}{\sigma(S_t)f_t}$$

We therefore have

$$\beta_T'(S_t)=\beta_S'(S_t)+\frac{\omega\rho}{f_t}\frac{d}{dS_t}\left(\frac{\nu(S_t,\sigma(S_t))}{\sigma(S_t)}\right)\tag{64}$$

$$=\beta_S'(S_t)+\frac{\omega\rho}{f_t}\left(\frac{\frac{d\nu(S_t,\sigma(S_t))}{dS_t}\sigma(S_t)-\frac{d\sigma(S_t)}{dS_t}\nu(S_t,\sigma(S_t))}{\sigma(S_t)^2}\right)$$

$$=\beta_S'(S_t)+\frac{\omega\rho}{f_t}\left(\frac{\left(\frac{\partial\nu(S_t,\sigma(S_t))}{\partial S_t}+\frac{\partial\nu(S_t,\sigma(S_t))}{\partial\sigma}\frac{d\sigma(S_t)}{dS_t}\right)\sigma(S_t)-\frac{d\sigma(S_t)}{dS_t}\nu(S_t,\sigma(S_t))}{\sigma(S_t)^2}\right)$$

Noting that $d\nu(S_t,\sigma(S_t))/dS_t$ is another representation of the the option's vanna, and that $d\nu(S_t,\sigma(S_t))/d\sigma$

is the option's *volga*, which is equal to $\nu(S_t,\sigma)d_1d_2/\sigma$, we have

$$\beta_T'(S_t) = \beta_S'(S_t) + \frac{\omega\rho}{f_t}\left(\frac{\left(\frac{\nu(S_t,\sigma)}{S_t}\left(1-\frac{d_1}{\sigma\sqrt{\tau}}\right) - \nu(S_t,\sigma(S_t))\frac{d_1d_2}{\sigma}\frac{\Delta(S_t,\sigma)}{\nu(S_t,\sigma)}\right)\sigma(S_t) + \frac{\Delta(S_t,\sigma)}{\nu(S_t,\sigma)}\nu(S_t,\sigma(S_t))}{\sigma(S_t)^2}\right)$$

$$= \beta_S'(S_t) + \frac{\omega\rho}{f_t}\left(\frac{\left(\frac{\nu(S_t,\sigma)}{S_t}\left(1-\frac{d_1}{\sigma\sqrt{\tau}}\right) - \Delta(S_t,\sigma(S_t))\frac{d_1d_2}{\sigma}\right)\sigma(S_t) + \Delta(S_t,\sigma)}{\sigma(S_t)^2}\right)$$

$$= \beta_S'(S_t) + \frac{\omega\rho}{\sigma(S_t)^2 f_t}\left(\frac{\nu(S_t,\sigma)\sigma(S_t)}{S_t}\left(1-\frac{d_1}{\sigma\sqrt{\tau}}\right) + \Delta(S_t,\sigma(S_t))\left(1-d_1d_2\right)\right)$$

Biases in relative price-weighted returns

In some cases we examine portfolios weighted according to the relative option price, or the price of the option divided by the price of the underlying stock. In this case our portfolio return is an equal weighted average of

$$w(\hat{S}_t)\left(\frac{\widehat{f}_{t+1}}{\widehat{f}_t} - 1\right) - w(\hat{S}_t)\beta(\hat{S})\left(\frac{S_{t+1}}{\hat{S}_t} - 1\right) \tag{65}$$

where

$$w(\hat{S}_t) \propto \frac{\hat{f}_t}{\hat{S}_t}$$

and the weights sum to N, the number of assets in the portfolio. We will assume that the constant of proportionality k is known exactly, which would approximately be the case in a large portfolio if measurement errors are independent.

The first component,

$$w(\hat{S}_t)\left(\frac{\widehat{f}_{t+1}}{\widehat{f}_t} - 1\right)$$

$$= \frac{k\hat{f}_t}{\hat{S}_t}\left(\frac{\widehat{f}_{t+1}}{\widehat{f}_t} - 1\right) = \frac{k}{\hat{S}_t}\left(\widehat{f}_{t+1} - \hat{f}_t\right) \tag{66}$$

is probably close to unbiased. The second component in the difference,

$$w(\hat{S}_t)\beta(\hat{S}_t)\left(\frac{S_{t+1}}{\hat{S}_t} - 1\right)$$

$$= \frac{k\hat{f}_t}{\hat{S}_t}\beta(\hat{S}_t)\left(\frac{S_{t+1}}{\hat{S}_t} - 1\right) = \beta(\hat{S}_t)\left(\frac{k\hat{f}_t S_{t+1}}{\hat{S}_t^2} - \frac{k\hat{f}_t}{\hat{S}_t}\right) \tag{67}$$

$$\equiv \beta(\hat{S}_t)h(\hat{S}_t)$$

We again determine the bias by taking half the second derivative and multiplying by the measurement error variance:

$$\frac{1}{2}\frac{d^2}{dS_t^2}\left[\beta(S_t)h(S_t)\right]\mathrm{Var}(S_t\delta)$$

$$= \frac{1}{2}\frac{d}{dS_t}\left[\beta'(S_t)h(S_t) + \beta(S_t)h'(S_t)\right]\mathrm{Var}(S_t\delta)$$

$$= \frac{1}{2}\left[\beta''(S_t)h(S_t) + \beta'(S_t)h'(S_t) + \beta'(S_t)h'(S_t) + \beta(S_t)h''(S_t)\right]\mathrm{Var}(S_t\delta)$$

$$= \frac{1}{2}\left[\beta''(S_t)h(S_t)S_t^2 + 2\beta'(S_t)h'(S_t)S_t^2 + \beta(S_t)h''(S_t)S_t^2\right]\mathrm{Var}(\delta)$$

Now note that

$$h(S_t)S_t^2 = \left(\frac{k\hat{f}_t S_{t+1}}{S_t^2} - \frac{k\hat{f}_t}{S_t}\right)S_t^2 = k\hat{f}_t S_{t+1} - k\hat{f}_t S_t$$

$$h'(S_t)S_t^2 = \left(-2\frac{k\hat{f}_t S_{t+1}}{S_t^3} + \frac{k\hat{f}_t}{S_t^2}\right)S_t^2 = -2\frac{k\hat{f}_t S_{t+1}}{S_t} + k\hat{f}_t$$

$$h''(S_t)S_t^2 = \left(6\frac{k\hat{f}_t S_{t+1}}{S_t^4} - 2\frac{k\hat{f}_t}{S_t^3}\right)S_t^2 = 6\frac{k\hat{f}_t S_{t+1}}{S_t^2} - 2\frac{k\hat{f}_t}{S_t}$$

Taking expectations over S_{t+1}, assuming it is a martingale, we have

$$\mathrm{E}\left[h(S_t)S_t^2\right] = 0$$

$$\mathrm{E}\left[h'(S_t)S_t^2\right] = -k\hat{f}_t$$

$$\mathrm{E}\left[h''(S_t)S_t^2\right] = 4\frac{k\hat{f}_t}{S_t}$$

The bias is therefore

$$\left[-k\hat{f}_t\beta'(S_t) + 2\frac{k\hat{f}_t}{S_t}\beta(S_t)\right]\mathrm{Var}(\delta)$$

References

[1] Ang, Andrew, Robert J. Hodrick, Yuhang Xing, and Xiaoyan Zhang, 2006, The Cross-Section of Volatility and Expected Returns, The Journal of Finance, 61, 259-299.

[2] Bakshi, Gurdip, Charles Cao, and Zhiwu Chen, 1997, Empirical Performance of Alternative Option Pricing Models, The Journal of Finance, 52, 2003-2049.

[3] Bakshi, Gurdip, Nikunj Kapadia and Dilip Madan, 2003, Stock Return Characteristics, Skew Laws, and Differential Pricing of Individual Equity Options, *Review of Financial Studies*, 17, 101-143.

[4] Bakshi, Gurdip and Nikunj Kapadia, 2003, Delta-Hedged Gains and the Negative Market Volatility Risk Premium, *Review of Financial Studies,* 16, 527-566.

[5] Bates, David S., 2000, Post-'87 Crash Fears in the S&P 500 Futures Option Market, *Journal of Econometrics,* 94, 181-238.

[6] Benzoni, Luca, 2002, Pricing Options under Stochastic Volatility: An Empirical Investigation, working paper, University of Minnesota.

[7] Black, Fischer and Myron Scholes, 1973, The Pricing of Options and Corporate Liabilities, Journal of Political Economy, 81, 637-654.

[8] Blume, M. and Robert Stambaugh, 1983, Bias in Computed Returns: An Application to the Size Effect, *Journal of Financial Economics,* 12, 387-553.

[9] Bollen, Nicolas P.B. and Robert E. Whaley, 2004, Does Net Buying Pressure Affect the Shape of Implied Volatility Functions?, *Journal of Finance,* 59, 711-753.

[10] Branger, Nicole and Christian Schlag, 2004, Can Tests Based on Option Hedging Errors Correctly Identify Volatility Risk Premia?, working paper. Goethe University.

[11] Britten-Jones, M. and A. Neuberger, 2000, Option Prices, Implied Price Processes, and Stochastic Volatility, The Journal of Finance, 55:2, 839-866.

[12] Chen, Joseph, 2003, Intertemporal CAPM and the Cross-Section of Stock Returns, working paper, University of Southern California.

[13] Chernov, Mikhail and Eric Ghysels, 2000, A Study Towards a Unified Approach to the Joint Estimation of Objective and Risk Neutral Measures for the Purpose of Options Valuation, Journal of Financial Economics, 56, 407-458.

[14] Christensen, B.J. and N.R. Prabhala, 1998, The Relation Between Implied and Realized Volatility, Journal of Financial Economics, 125-150.

[15] Coval, Joshua D. and Tyler Shumway, 2001, Expected Option Returns, *Journal of Finance,* 56, 983-1009.

[16] Driessen, Joost, Pascal Maenhout and Grigory Vilkov, 2006, Option-Implied Correlations and the Price of Correlation Risk, working paper, University of Amsterdam.

[17] Duarte, Jefferson and Lance Young, 2007, Why is PIN priced?, working paper. University of Washington.

[18] Duarte, Jefferson, 2004, Evaluating An Alternative Risk Preference in Affine Term Structure Models Review of Financial Studies vol. 17, no. 2, pp.370-404.

[19] Dubinsky, Andrew and Michael Johannes, 2005, Earnings Announcements and Equity Options, working paper, Columbia University.

[20] Fleming, Jeff, Barbara Ostdiek and Robert Whaley, 1995, Predicting Stock Market Volatility: A New Measure, Journal of Futures Markets, 15:3, 265-302.

[21] 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.

[22] Jackwerth, Jens Carsten and Mark Rubinstein, 1996, Recovering Probability Distributions from Option Prices, The Journal of Finance, 51, 1611-1631.

[23] Jones, Christopher S., 2003, The Dynamics of Stochastic Volatility: Evidence From Underlying and Options Markets, *Journal of Econometrics,* 116, 181-224.

[24] Lee, Charles M.C. and Mark J. Ready, 1991, Inferring trade direction from intraday data, *Journal of Finance* 46, 733-746.

[25] Merton, R., 1973, An Intertemporal Capital Asset Pricing Model, Econometrics, 41, 867-888.

[26] Ni, Sophie X., 2006, Stock Option Returns: A Puzzle, Working paper University of Illinois at Urbana-Champaign

[27] Pan, Jun, 2002, The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study, Journal of Financial Economics, 63, 3-50.

[28] Santa-Clara, Pedro and Alessio Saretto, 2005, Option Strategies: Good Deals and Margin Calls, working paper, UCLA.

Table 1 - Summary statistics. The first panel of this table contains the distributions of some variables of interest. The implied volatility (σ) is from the volatility surface file of the options database and it is the implied volatility of a hypothetical option with time-to-maturity equal to 30 days and delta equal to 0.5. The correlation between stock returns and changes in the implied volatility is represented by ρ. The annualized volatility of volatility (ω) is estimated from the daily changes in at-the-money volatility. The betas with respect to market (β_M) and the market volatility factor (β_{VOL}) are the pre-ranking betas used to form portfolios for the Fama-MacBeth regressions.

	Mean	Standard Deviation	1st Percentile	25th Percentile	50th Percentile	75th Percentile	99th Percentile
σ	0.51	0.28	0.20	0.31	0.45	0.65	1.04
ρ	-0.25	0.15	-0.47	-0.34	-0.26	-0.17	-0.03
ω	1.04	0.90	0.30	0.52	0.84	1.27	2.40
β_m	-0.08	0.70	-0.77	-0.24	-0.06	0.10	0.55
β_{vol}	0.34	1.86	-0.71	-0.03	0.23	0.58	1.63

	Daily delta-hedged returns						
	Mean	Standard Deviation	1st Percentile	25th Percentile	50th Percentile	75th Percentile	99th Percentile
Delta-Hedge	2200	3061	1112	1450	1734	2145	4104
Total-Delta minus Delta-Hedged	30	822	-187	-44	-5	25	187

Table 2 - Double sorts with the actual data. This table displays mean daily returns as well as mean absolute and relative bid-ask spreads of options with different moneyness and maturities. Mean returns are in basis points. Moneyness is measured in terms of the number of standard deviations the option is from being at-the-money and ranges from –3 to +3. There are three time-to-maturity groups. Short-term options expire within 10-30 days, medium term options expire within 31-120 days, and long-term options expire within 121-260 days. T-statistics are in parentheses.

Mean daily option returns censoring options without implied volatility

Calls				Puts			
Moneyness	Short	Medium	Long	Moneyness	Short	Medium	Long
Deep ITM	-30 (-44.48)	-29 (-88.52)	-51 (-74.30)	Deep OTM	20 (4.05)	83 (25.56)	31 (7.57)
	-1 (-2.16)	-2 (-11.32)	-7 (-36.28)		345 (103.36)	162 (120.57)	61 (43.11)
	25 (36.48)	18 (86.68)	8 (49.61)		77 (32.97)	33 (46.49)	17 (30.47)
ATM	46 (58.21)	39 (169.24)	26 (175.34)	ATM	-12 (-14.04)	-3 (-9.97)	2 (14.50)
	213 (99.42)	110 (154.14)	66 (124.12)		-10 (-11.42)	-11 (-42.32)	-2 (-10.37)
	557 (195.08)	312 (262.41)	183 (151.90)		-22 (-37.41)	-17 (-85.13)	-5 (-32.38)
Deep OTM	271 (65.29)	326 (133.81)	262 (92.60)	Deep ITM	-34 (-51.72)	-25 (-104.89)	-8 (-36.67)

Mean daily option returns not censoring options without implied volatility

Calls				Puts			
Moneyness	Short	Medium	Long	Moneyness	Short	Medium	Long
Deep ITM	13 (33.86)	0 (-0.22)	-15 (-39.61)	Deep OTM	307 (47.50)	203 (62.87)	130 (33.71)
	21 (51.83)	9 (56.47)	0 (-2.31)		406 (128.00)	181 (130.77)	70 (49.38)
	34 (48.21)	20 (94.69)	9 (56.11)		98 (39.44)	34 (45.91)	16 (26.02)
ATM	49 (59.53)	38 (161.52)	25 (166.41)	ATM	-9 (-9.49)	-4 (-13.36)	1 (8.92)
	236 (100.32)	111 (144.78)	64 (121.01)		-1 (-1.71)	-9 (-33.66)	-1 (-6.29)
	664 (237.60)	334 (276.84)	186 (151.71)		-1 (-1.62)	-9 (-48.49)	-2 (-12.26)
Deep OTM	574 (147.79)	427 (183.30)	308 (109.94)	Deep ITM	-5 (-12.32)	-9 (-53.05)	-1 (-6.25)

Mean option relative bid-ask spreads not censoring options without implied volatility

Calls				Puts			
Moneyness	Short	Medium	Long	Moneyness	Short	Medium	Long
Deep ITM	0.06	0.04	0.03	Deep OTM	1.72	1.60	1.36
	0.08	0.05	0.04		1.11	0.91	0.64
	0.11	0.07	0.06		0.48	0.32	0.23
ATM	0.2	0.13	0.10	ATM	0.21	0.14	0.10
	0.49	0.3	0.20		0.11	0.08	0.06
	1.12	0.78	0.52		0.08	0.06	0.04
Deep OTM	1.75	1.54	1.21	Deep ITM	0.06	0.05	0.03

Mean option absolute bid-ask spreads not censoring options without implied volatility

Calls				Puts			
Moneyness	Short	Medium	Long	Moneyness	Short	Medium	Long
Deep ITM	0.45	0.47	0.44	Deep OTM	0.18	0.17	0.15
	0.39	0.46	0.45		0.19	0.18	0.17
	0.33	0.41	0.44		0.20	0.21	0.21
ATM	0.25	0.30	0.34	ATM	0.25	0.28	0.31
	0.19	0.21	0.22		0.31	0.36	0.39
	0.18	0.19	0.18		0.36	0.41	0.41
Deep OTM	0.18	0.18	0.16	Deep ITM	0.40	0.44	0.42

Table 3 - Parameters used in the simulations. The first panel of this table displays the mean (M) of the logs of the option bid-ask spreads used in the simulations. The second panel displays the standard deviations (S) of the logs of options bid-spreads in the simulations. The bid-ask spread of an option in the simulation is set equal to exp(M+S μ_i), where M and S are given by the displayed functions of the time-to-maturity and moneyness of the option and μ_i is a normally-distributed random number common to all options on a given underlying stock. The third and fourth panels display additional parameters used in the simulations. U[x,y] represents the uniform distribution between x and y.

Mean of log spread (M)

Calls				Puts			
Moneyness	Short	Medium	Long	Moneyness	Short	Medium	Long
Deep ITM	-3.69	-3.89	-3.93	Deep OTM	-0.45	-0.49	-0.43
	-3.32	-3.56	-3.48		-1.16	-1.34	-1.17
	-2.77	-3.12	-2.96		-1.89	-2.23	-2.03
ATM	-2.23	-2.59	-2.47	ATM	-2.44	-2.76	-2.60
	-1.55	-1.92	-1.79		-3.01	-3.31	-3.18
	-0.59	-0.77	-0.49		-3.61	-3.81	-3.77
Deep OTM	0.19	0.25	0.37	Deep ITM	-4.03	-4.25	-4.35

Standard deviation of log spread (S)

Calls				Puts			
Moneyness	Short	Medium	Long	Moneyness	Short	Medium	Long
Deep ITM	0.90	0.91	0.92	Deep ITM	1.04	1.09	1.09
	0.88	0.88	0.85		1.23	1.22	1.16
	0.86	0.87	0.77		1.13	1.07	0.96
ATM	0.88	0.88	0.73	ATM	0.96	0.92	0.83
	1.07	1.05	0.91		0.90	0.88	0.87
	1.15	1.17	1.01		0.88	0.89	0.92
Deep OTM	0.62	0.62	0.58	Deep OTM	0.89	0.91	0.93

Market-level parameters

$r_t - q_t$	0.04/252	risk-neutral drift
λ_1	0.0125	price of market risk
λ_2	0 or -0.1	price of volatility risk
κ_V^m	0.018	variance mean-reversion parameter
$\overline{V}^m$	0.00013	long-run risk-neutral mean of variance
ω_V^m	0.0028	volatility of variance parameter
ρ^m	-0.7	correlation between return and change in variance

Firm-level parameters

κ_V^i	U[0.01,0.05]	variance mean reversion parameter
$\overline{V}^i$	U[0.0002,0.001]	long-run risk neutral mean of variance
ω_V^i	U[0.002,0.005]	volatility of variance parameter
ρ^i	U[-0.5,-0.1]	correlation between return and change in variance
ξ_1^i	U[0.25,0.75]	correlation between market and equity returns
ξ_2^i	U[0.25,0.75]	correlation between market and equity changes in variance
error sd	0 or U[0.001,0.005]	standard deviation of stock price measurement error

Table 4 - Double sorts with simulated data. This table displays mean daily excess returns from simulated option prices. The simulated model is a parameterization of the Heston model, where the price of volatility risk is negative. The first panel contains the mean returns of simulated option prices without any bid-ask bounce. The second panel contains the mean returns of simulated option prices that are subjected to bid-ask bounce but that are otherwise identical. The third panel contains mean returns, also computed with bid-ask bounce, where options are censored if their prices do not satisfy arbitrage bounds. Moneyness and maturity are as defined in Table 2. Mean returns are in basis points, and t-statistics are in parentheses.

Mean of simulated option returns with no errors

Calls				Puts			
Moneyness	Short	Medium	Long	Moneyness	Short	Medium	Long
Deep ITM	6 (11.6)	3 (6.9)	3 (5.8)	Deep ITM	-157 (62.3)	-105 (35.1)	-82 (27.0)
	6 (14.3)	3 (9.9)	3 (8.7)		-122 (50.9)	-78 (31.0)	-58 (23.9)
	1 (21.3)	0 (14.4)	1 (12.3)		-77 (38.8)	-50 (24.4)	-39 (19.7)
ATM	-10 (28.6)	-8 (19.2)	-5 (16.3)	ATM	-48 (29.9)	-32 (19.2)	-25 (15.7)
	-36 (38.2)	-25 (25.4)	-17 (21.2)		-27 (21.9)	-18 (14.4)	-14 (11.7)
	-87 (53.0)	-61 (34.2)	-44 (28.1)		-12 (14.1)	-7 (9.0)	-6 (7.3)
Deep OTM	-104 (72.5)	-91 (37.7)	-70 (27.8)	Deep OTM	-6 (11.0)	-4 (5.9)	-3 (4.6)

Mean of simulated option returns with errors and no censoring

Calls				Puts			
Moneyness	Short	Medium	Long	Moneyness	Short	Medium	Long
Deep ITM	8 (11.3)	4 (6.8)	4 (5.8)	Deep ITM	570 (76.2)	580 (52.2)	659 (48.8)
	8 (14.3)	5 (9.7)	5 (8.6)		294 (56.6)	305 (39.3)	314 (34.5)
	10 (21.0)	5 (14.0)	5 (12.2)		66 (40.5)	43 (25.8)	50 (21.6)
ATM	27 (28.1)	11 (18.9)	9 (16.1)	ATM	-18 (30.0)	-17 (19.1)	-12 (15.7)
	187 (40.5)	126 (27.1)	123 (24.3)		-20 (21.7)	-14 (14.0)	-11 (11.6)
	602 (63.7)	597 (46.3)	686 (45.1)		-10 (14.0)	-6 (8.8)	-5 (7.3)
Deep OTM	1255 (117.5)	1144 (61.6)	1336 (58.5)	Deep OTM	-5 (10.7)	-3 (5.8)	-3 (4.7)

Mean of simulated option returns with errors and censoring

Calls				Puts			
Moneyness	Short	Medium	Long	Moneyness	Short	Medium	Long
Deep ITM	-27 (12.7)	-19 (7.5)	-18 (6.5)	Deep ITM	327 (61.4)	437 (47.5)	518 (50.1)
	-10 (15.2)	-6 (10.1)	-4 (8.9)		237 (44.3)	277 (32.9)	297 (31.0)
	5 (21.3)	3 (14.2)	4 (12.3)		63 (39.1)	42 (25.4)	49 (21.4)
ATM	26 (28.2)	10 (18.9)	9 (16.1)	ATM	-18 (29.9)	-17 (19.1)	-12 (15.7)
	181 (42.2)	125 (27.6)	123 (24.5)		-23 (21.5)	-15 (13.9)	-12 (11.4)
	492 (87.3)	541 (60.2)	644 (56.0)		-24 (13.9)	-15 (8.6)	-14 (7.1)
Deep OTM	691 (189.7)	860 (121.8)	1145 (99.7)	Deep OTM	-44 (10.6)	-23 (5.5)	-18 (4.5)

Table 5 - Fama-MacBeth regressions on simulated data. This table displays the means and the standard deviations (in parentheses) of the estimated coefficients of Fama-MacBeth regressions and their t-statistics. The means and standard deviations are calculated across all simulations. The simulated model is a parameterization of the Heston model with negative volatility risk premium. The first panel displays the results based on simulated data without bid-ask bounce and without any filtering of options. The second panel displays the results based on simulated data with bid-ask bounce. The third panel displays the results based on samples that have bid-ask bounce that filter out options that have large relative bid-ask spreads prior to the start of the interval over which returns are computed. The regressions are estimated with portfolios of total delta-hedged options sorted by time-to-maturity, moneyness, and pre-ranking volatility beta. The variables in the cross-sectional regressions are the betas with the respect to the market (β_M), the beta with respect to the market volatility (β_{VOL}), the mean gamma (γ), the mean of the square of the option spreads (OPTION SPREAD) in the portfolio and the mean of the bias adjustment for stock spreads ($\beta\times$R BIAS),), which is given computed as $[b_T^{'}(S_t)S_t + b_T(S_t)]'$ Stock Spread^2, where $b_T(S_t)$ is the Black-Scholes beta of the option with respect to the stock.

No errors, no filtering

	1		2		3	
	coefficient	t-statistic	coefficient	t-statistic	coefficient	t-statistic
Intercept x 10^4	0.35	0.47	-0.01	0.02	0.07	0.16
	(0.83)	(1.06)	(0.65)	(1.03)	(0.6)	(1.05)
β_M x 10^3			0.19	0.69	0.17	0.60
			(0.30)	(1.04)	(0.313)	(1.05)
β_{Vol} x 10^2	-0.005	-4.79	-0.005	-4.7	-0.005	-4.64
	(0.001)	(1.04)	(0.001)	(1.05)	(0.001)	(1.01)
γ x 10^6					-0.22	-0.14
					(2.093)	(1.13)

Option errors, no filtering

	1		2		3		4	
	coefficient	t-statistic	coefficient	t-statistic	coefficient	t-statistic	coefficient	t-statistic
Intercept x 10^4	-38.48	-18.01	-47.63	-21.37	-19.86	-5.63	-24.36	-5.42
	(6.63)	(4.71)	(8.28)	(6.82)	(4.79)	(1.11)	(6.14)	(1.39)
β_M x 10^3			9.45	11.50			1.78	1.88
			(5.505)	(7.21)			(1.401)	(1.42)
β_{Vol} x 10^2	0.036	28.96	0.039	30.41	-0.009	-6.29	-0.009	-5.84
	(0.004)	(2.93)	(0.005)	(3.92)	(0.002)	(1.23)	(0.002)	(1.26)
OPTION					0.08	44.58	0.08	45.30
SPREAD x 10^4					(0.003)	(9.49)	(0.002)	(8.71)
					6.28	6.81	7.09	6.45
$\beta\times$R BIAS					(1.382)	(1.38)	(1.517)	(1.39)

Option errors and filtering on t-2 spread

	1		2	
	coefficient	t-statistic	coefficient	t-statistic
Intercept x 10^4	8.97	10.18	-2.90	-1.39
	(1.13)	(1.64)	(2.09)	(0.98)
β_M x 10^3	-0.02	-0.03	0.69	1.87
	(0.336)	(0.99)	(0.37)	(1.00)
β_{Vol} x 10^2	-0.003	-2.6	-0.005	-4.13
	(0.001)	(1)	(0.001)	(1.07)
OPTION			0.04	3.28
SPREAD x 10^4			(0.01)	(0.98)
			2.89	4.51
$\beta\times$R BIAS			(0.686)	(1.11)

Table 6 - Fama-MacBeth regressions on simulated data with zero volatility risk premium. This table displays the means and the standard deviations (in parentheses) of the estimated coefficients of Fama-MacBeth regressions and their t-statistics. The means and standard deviations are calculated across all simulations. The simulated model is a parameterization of the Heston model with zero volatility risk premium. All regressions are based on simulated data with bid-ask bounce. The regressions with filtering are based on samples with deletion of options with large relative bid-ask spreads prior to the start of the return interval. The regressions are estimated with portfolios of total delta-hedged options sorted by time-to-maturity, moneyness, and pre-ranking volatility beta. The variables in the cross-sectional regressions are the beta with respect to the market (β_M), the beta with respect to the market volatility (β_{VOL}), the mean of the square of the option spreads (OPTION SPREAD) in the portfolio, and the mean of the bias adjustment for stock spreads ($\beta\times$R BIAS), which is given by $[b_T^{'}(S_t)S_t + b_T(S_t)]$´ Stock Spread^2, where $b_T(S_t)$ is the Black and Scholes beta of the option with respect to the stock.

	No filtering				Filtering on t-2 spread	
	1		2		3	
	Coefficient	t-statistic	Coefficient	t-statistic	Coefficient	t-statistic
Intercept x 10^4	-45.20	-21.67	-20.94	-4.93	-2.92	-1.44
	(7.79)	(6.40)	(5.62)	(1.20)	(2.10)	(1.02)
β_M x 10^3	14.49	11.91	1.73	1.17	0.89	1.69
	(7.49)	(6.58)	(2.02)	(1.17)	(0.51)	(0.97)
β_{Vol} x 10^2	0.050	32.57	-0.004	-2.42	0.000	0.06
	(0.006)	(3.69)	(0.002)	(1.12)	(0.001)	(1.01)
OPTION SPREAD x 10^4			0.08	45.84	0.04	3.43
			(0.002)	(8.56)	(0.011)	(0.98)
$\beta\times$R BIAS			6.92	5.76	2.89	3.96
			(1.61)	(1.29)	(0.79)	(1.10)

Table 7 – S&P 500 options returns. The first panel displays the results of a Fama-MacBeth regression on S&P 500 option returns. The variables in the cross-sectional regressions are the beta with the respect to the market (β_M), the beta with respect to the market volatility (β_{VOL}) and the mean of the square of the option spreads (OPTION SPREAD) in the portfolio. The second panel contains the average returns of S&P 500 options across moneyness and time-to-maturity. The third panel displays the spread bias component of the average returns. This spread bias component is calculated with the coefficient on the OPTION SPREAD in the first panel and the bid-ask spread of each option in the sample. The third panel displays the difference between the average return in the second panel and the bias component, this difference is the bias adjusted average option return.

Option-spread-bias-correction Fama-MacBeth regression

Intercept	β_{Vol}	β_M	OPTION SPREAD	Average number of cross-sectional observations	Number of time series observations
0.0007	-0.0003	-0.00004	0.04369	62	2402
(2.2)	(-4.6)	(-0.2)	(4.3)		

Average returns

Calls				Puts			
Moneyness	Short	Medium	Long	Moneyness	Short	Medium	Long
Deep ITM	10.9	11.9	-5.2	Deep OTM	-358.4	-81.3	-20.0
	10.6	25.1	-1.9		-260.9	-56.7	-32.2
	3.5	35.3	12.2		-139.4	-44.8	-17.7
ATM	-6.5	38.7	8.5	ATM	-76.9	-34.1	-13.8
	-15.7	57.5	22.1		-42.9	-48.0	-12.6
	27.6	190.5	46.7		-25.4	-12.4	0.1
Deep OTM	830.2	426.7	237.8	Deep ITM	-25.6	9.4	15.0

Spread bias component

Calls				Puts			
Moneyness	Short	Medium	Long	Moneyness	Short	Medium	Long
Deep ITM	0.1	0.0	0.0	Deep OTM	43.9	20.9	12.5
	0.3	0.1	0.0		12.7	6.7	4.5
	0.9	0.3	0.1		4.6	2.9	1.5
ATM	2.1	1.0	0.4	ATM	2.2	1.1	0.5
	7.4	4.1	2.1		1.1	0.4	0.2
	183.8	142.2	76.2		0.3	0.1	0.0
Deep OTM	1139.0	1413.8	1141.3	Deep ITM	0.1	0.0	0.0

Bias-adjusted average returns

Calls				Puts			
Moneyness	Short	Medium	Long	Moneyness	Short	Medium	Long
Deep ITM	10.8	11.9	-5.2	Deep OTM	-402.3	-102.2	-32.5
	10.3	25.1	-1.9		-273.6	-63.4	-36.6
	2.6	35.0	12.1		-144.0	-47.7	-19.2
ATM	-8.6	37.7	8.2	ATM	-79.1	-35.3	-14.2
	-23.1	53.3	20.0		-44.0	-48.4	-12.8
	-156.2	48.2	-29.5		-25.8	-12.5	0.1
Deep OTM	-308.8	-987.1	-903.5	Deep ITM	-25.7	9.3	15.0

Table 8 - Fama-MacBeth regressions on portfolios of equity options. This table displays the results of Fama-MacBeth regressions on 420 portfolios formed on the basis of moneyness, maturity, and pre-ranking volatility beta. Calls and puts comprise separate portfolios. The sample period is 1997 to 2006. Options with bid-ask spread larger than 25% of the mid option price are deleted. The dependent variables are total delta-hedged returns. The variables included in in the cross-sectional regressions are the beta with the respect to market volatility (β_{Vol}), the beta with respect to the stock market return (β_M), the mean of the square of the option's relative spread (OPTION SPREAD), the mean of $b_T^{'}(S_t)S_t + b_T(S_t)$ where $b_T(S_t)$ is the Black and Scholes beta of the option with respect to the stock ($\beta\times$R BIAS INTERCEPT), and the mean of $[b_T^{'}(S_t)S_t + b_T(S_t)]' \text{ Stock Spread}^2$ ($\beta\times$R BIAS SLOPE). The variables $\beta\times$R BIAS SLOPE and $\beta\times$R BIAS INTERCEPT are adjustments for the bias in the delta hedged return caused by stock bid-ask bounce. RMSE is the root mean square error of the cross-sectional regressions. Newey-West t-statistics, computed using an automatic lag length selector, are in parentheses.

	Tripled sorted portfolios with filters					
	Total delta-hedged. EW portfolios, OLS regressions					
	1	2	3	4	5	6
Intercept $\times 10^3$	0.176	0.877	0.322	0.186	0.038	0.194
	(1.015)	(1.844)	(2.054)	(1.145)	(0.137)	(1.414)
$\beta_{Vol} \times 10^3$	0.328		0.224	-0.193		-0.255
	(2.400)		(1.506)	(-1.659)		(-2.002)
$\beta_M \times 10^3$		-1.603	-0.992		0.601	0.132
		(-3.294)	(-2.887)		(1.786)	(0.378)
OPTION SPREAD				0.032	0.025	0.036
				(9.606)	(3.597)	(10.171)
$\beta\times$R BIAS				1.364	0.941	-0.444
INTERCEPT $\times 10^7$				(1.175)	(0.833)	(-0.408)
				-0.073	-0.097	-0.080
$\beta\times$R BIAS SLOPE				(-1.280)	(-1.878)	(-1.553)
R-square	0.530	0.312	0.590	0.651	0.564	0.567
RMSE $\times 10^4$	13.0	15.8	12.2	11.2	12.6	12.5
Avg. CS obs.	266.7	266.7	266.7	266.7	266.7	266.7
TS obs.	2240	2240	2240	2240	2240	2240

Table 9 – Robustness analysis for Fama-MacBeth regressions on portfolios of equity options. This table displays the results of Fama-MacBeth regressions on 420 portfolios formed on the basis of moneyness, maturity, and pre-ranking volatility beta. Calls and puts comprise separate portfolios. The sample period is 1997 to 2006. Options with bid-ask spread larger than 25% of the mid option price are deleted. The dependent variables are total delta-hedged returns. The variables in the cross-sectional regressions are the beta with the respect to market volatility (β_{Vol}), the beta with respect to the stock market return (β_M), and three bias control variables that were described in Table 8. RMSE is the root mean square error of the cross-sectional regressions. Newey-West t-statistics, computed using an automatic lag length selector, are in parentheses.

	Tripled sorted portfolios with filters			
	OLS Regression		WLS Regression	
	Equal-Weighted	Relative Price-Weighted	Equal-Weighted	Relative Price-Weighted
Intercept $\times 10^3$	0.194	0.433	-0.169	0.024
	(1.414)	(3.275)	(-7.770)	(1.024)
$\beta_{Vol} \times 10^3$	-0.255	-0.480	0.288	0.170
	(-2.002)	(-3.862)	(1.868)	(1.164)
$\beta_M \times 10^3$	0.132	0.026	1.312	1.248
	(0.378)	(0.068)	(3.807)	(3.475)
OPTION SPREAD	0.036	0.007	0.022	-0.012
	(10.171)	(1.648)	(5.814)	(-2.969)
$\beta \times$R BIAS	-0.444	-45.306	0.138	-8.572
INTERCEPT $\times 10^7$	(-0.408)	(-5.542)	(0.410)	(-2.979)
	-0.080	-0.051	0.004	-0.009
$\beta \times$R BIAS SLOPE	(-1.553)	(-0.889)	(0.318)	(-0.529)
R-square	0.567	0.532	0.118	0.037
RMSE $\times 10^4$	12.5	12.5	17.9	17.9
Avg. CS obs.	266.7	266.7	266.7	266.7
TS obs.	2240	2240	2240	2240

Table 10 – The predictive ability of volatility risk premia. This table measures the forecastability of risk premia and factors. The explanatory variable in the top panel is the realized price of volatility risk, lagged one day, measured as the coefficient on volatility betas in a single cross-sectional regression. The contemporaneous version of this variable is the first dependent variable. The realized price of market risk is the second dependent variable. The volatility and market return factors are the last two dependent variables. The second panel is identical except that the explanatory variable is a 22-day moving average of the realized price of volatility risk. White T-statistics are in parentheses.

Dependent variable:	λ_{Vol}	λ_M	volatility factor	stock return factor
Intercept $\times 10^3$	-0.213	0.246	0.077	0.335
	-(2.539)	(0.678)	(3.883)	(1.309)
1-day lag of λ_{Vol}	0.165	0.452	-0.012	0.131
	(6.012)	(1.755)	-(1.664)	(1.273)
R-square	0.027	0.011	0.002	0.001
# of obs.	2239	2239	2239	2239

Dependent variable:	λ_{Vol}	λ_M	volatility factor	stock return factor
Intercept $\times 10^3$	-0.140	0.449	0.085	0.405
	-(1.585)	(1.223)	(3.840)	(1.457)
lagged 1-month	0.464	1.286	0.018	0.508
average of λ_{Vol}	(6.247)	(3.809)	(0.888)	(1.991)
R-square	0.022	0.010	0.000	0.003
# of obs.	2218	2218	2218	2218

Table 11 - Fama-MacBeth regressions with time varying risk premia. This table displays the results of Fama-MacBeth regressions on 420 portfolios formed on the basis of moneyness, maturity, and pre-ranking volatility beta. Calls and puts comprise separate portfolios. The sample period is 1997 to 2006. Options with bid-ask spread larger than 25% of the mid option price are deleted. The dependent variables are total delta-hedged returns. The variables in the cross-sectional regressions are the beta with the respect to market volatility (β_{Vol}), the beta with respect to the stock market return (β_M), and three variables described in Table 8 that control for bid-ask bounce. Cross-sectional regression intercepts and slope coefficients on β_{Vol} and β_M are regressed on the lag of the VIX index, while only unconditional means are reported for the bias controls. Newey-West t-statistics, computed using an automatic lag length selector, are in parentheses.

	Total delta-hedged. EW portfolios, OLS regressions			
	1	2	3	4
Intercept $\times 10^3$	0.322	-2.110	0.194	-1.869
	(2.054)	-(3.034)	(1.414)	-(3.208)
Intercept $\times$ VIX(t-1) $\times 10^2$		1.101		0.934
		(3.362)		(3.400)
$\beta_{Vol} \times 10^3$	0.022	-1.652	-0.026	-2.032
	(1.506)	-(3.390)	-(2.002)	-(4.541)
$\beta_{Vol} \times$ VIX(t-1) $\times 10^2$		0.849		0.804
		(3.522)		(3.754)
$\beta_M \times 10^3$	-0.992	-2.801	0.132	-4.467
	-(2.887)	-(1.878)	(0.378)	-(2.948)
$\beta_M \times$ VIX(t-1) $\times 10^2$		0.819		2.082
		(1.223)		(2.990)
OPTION SPREAD			0.036	0.036
			(10.171)	(10.171)
$\beta \times$R BIAS INTERCEPT $\times 10^7$			-0.444	-0.444
			-(0.408)	-(0.408)
			-0.080	-0.080
$\beta \times$R BIAS SLOPE			-(1.553)	-(1.553)
R-square	0.590	0.586	0.567	0.572
RMSE $\times 10^4$	12.2	12.2	12.5	12.4
Avg. CS obs.	266.7	266.7	266.7	266.7
TS obs.	2240	2240	2240	2240

Table 12 – Robustness analysis for Fama-MacBeth regressions with time varying risk premia. This table displays the results of Fama-MacBeth regressions on 420 portfolios formed on the basis of moneyness, maturity, and either the pre-ranking volatility beta, the stock's at-the-money implied volatility, or the ratio of implied to 22-day trailing realized volatility. Calls and puts comprise separate portfolios. The sample period is 1997 to 2006. Options with bid-ask spread larger than 25% of the mid option price are deleted. The dependent variables are either total or regular delta-hedged returns. The variables in the cross-sectional regressions are the beta with the respect to market volatility (β_{Vol}), the beta with respect to the stock market return (β_M), and three variables described in Table 8 that control for bid-ask bounce. Cross-sectional regression intercepts and slope coefficients on β_{Vol} and β_M are regressed on the lag of the VIX index, while only unconditional means are reported for the bias controls. Newey-West t-statistics, computed using an automatic lag length selector, are in parentheses.

	Total delta-hedged						Regular delta-hedged
Sort variable:	β_{Vol}	β_{Vol}	β_{Vol}	β_{Vol}	Implied Volatility	Volatility Ratio	β_{Vol}
	OLS EW	OLS RPW	WLS EW	WLS RPW	OLS EW	OLS EW	OLS EW
Intercept × 10^3	-1.869	-1.123	-0.381	-0.197	-1.589	-1.652	-0.486
	-(3.208)	-(1.841)	-(5.332)	-(1.794)	-(3.177)	-(2.999)	-(5.611)
Intercept × VIX(t-1) × 10^2	0.934	0.704	0.096	0.100	0.796	0.745	0.139
	(3.400)	(2.496)	(2.726)	(1.841)	(3.462)	(2.902)	(3.210)
β_{Vol} × 10^3	-2.032	-1.778	-1.942	-1.690	-1.834	-2.044	-2.140
	-(4.541)	-(3.985)	-(3.506)	-(3.227)	-(4.094)	-(4.953)	-(3.796)
β_{Vol} × VIX(t-1) × 10^2	0.804	0.588	1.009	0.842	0.723	0.922	1.064
	(3.754)	(2.787)	(3.696)	(3.318)	(3.285)	(4.401)	(3.812)
β_M × 10^3	-4.467	-5.663	-4.068	-3.336	-5.795	-5.458	-3.391
	-(2.948)	-(3.355)	-(2.801)	-(2.139)	-(3.707)	-(3.358)	-(2.305)
β_M × VIX(t-1) × 10^2	2.082	2.575	2.435	2.075	2.768	2.463	1.439
	(2.990)	(3.475)	(3.513)	(2.870)	(3.828)	(3.249)	(2.135)
OPTION SPREAD	0.036	0.007	0.022	-0.012	0.037	0.019	0.023
	(10.171)	(1.648)	(5.814)	-(2.969)	(11.032)	(5.950)	(5.665)
β×R BIAS INTERCEPT × 10^7	-0.444	-45.306	0.138	-8.572	0.558	-0.403	-0.209
	-(0.408)	-(5.542)	(0.410)	-(2.979)	(0.393)	-(0.382)	-(0.681)
β×R BIAS SLOPE	-0.080	-0.051	0.004	-0.009	-0.018	0.032	-0.006
	-(1.553)	-(0.889)	(0.318)	-(0.529)	-(0.387)	(0.710)	-(0.441)
R-square	0.572	0.539	0.133	0.053	0.464	0.300	0.307
RMSE × 10^4	12.4	12.4	17.7	17.8	16.2	21.1	15.9
Avg. CS obs.	266.7	266.7	266.7	266.7	263.5	263.6	266.7
TS obs.	2240	2240	2240	2240	2240	2240	2240

Table 13 – Decomposing volatility betas into contract and firm components. This table displays the results of Fama-MacBeth regressions on 420 portfolios formed on the basis of moneyness, maturity, and pre-ranking volatility beta. Calls and puts comprise separate portfolios. The sample period is 1997 to 2006. Options with bid-ask spread larger than 25% of the mid option price are deleted. The dependent variables are total delta-hedged returns. In specifications 1 and 3, the variables in the cross-sectional regressions are the beta with the respect to market volatility (β_{Vol}), the beta with respect to the stock market return (β_M), and three variables described in Table 8 that control for bid-ask bounce. Specifications differ by separating β_{Vol} into "contract" and "firm" components. The contract component is identical for all portfolios with the same maturity and moneyness. The firm component, which accounts for variation due to differences in the underlying stocks, is the difference between the portfolio's own estimated beta and the contract component. In specifications 1 and 2, risk premia are assumed constant. In 3 and 4, cross-sectional regression intercepts and slope coefficients on β_{Vol} and β_M are regressed on the lag of the VIX index. Newey-West t-statistics, computed using an automatic lag length selector, are in parentheses.

	1	2	3	4
Intercept $\times 10^3$	0.194	0.312	-1.869	-1.668
	(1.414)	(2.423)	(-3.208)	(-3.146)
Intercept $\times$ VIX(t-1) $\times 10^2$			0.934	0.896
			(3.400)	(3.622)
Total $\beta_{Vol} \times 10^3$	-0.255		-2.032	
	(-2.002)		(-4.541)	
Contract component of $\beta_{Vol} \times 10^3$		-0.402		-2.264
		(-2.766)		(-4.378)
Firm component of $\beta_{Vol} \times 10^3$		0.253		-1.226
		(1.946)		(-2.852)
Total $\beta_{Vol} \times$ VIX(t-1) $\times 10^2$			0.804	
			(3.754)	
Contract component of $\beta_{Vol} \times$ VIX(t-1) $\times 10^2$				0.843
				(3.429)
Firm component of $\beta_{Vol} \times$ VIX(t-1) $\times 10^2$				0.669
				(3.533)
$\beta_M \times 10^3$	0.132	0.132	-4.467	-4.439
	(0.378)	(0.381)	(-2.948)	(-2.914)
$\beta_M \times$ VIX(t-1) $\times 10^2$			2.082	2.069
			(2.990)	(2.953)
OPTION SPREAD	0.036	0.043	0.036	0.043
	(10.171)	(14.658)	(10.171)	(14.658)
	-0.444	-0.964	-0.444	-0.964
$\beta \times$R BIAS INTERCEPT $\times 10^7$	(-0.408)	(-0.859)	(-0.408)	(-0.859)
	-0.080	-0.061	-0.080	-0.061
$\beta \times$R BIAS SLOPE	(-1.553)	(-1.206)	(-1.553)	(-1.206)
R-square	0.567	0.561	0.572	0.566
RMSE $\times 10^4$	12.5	12.6	12.4	12.5
Avg. CS obs.	266.7	266.7	266.7	266.7
TS obs.	2240	2240	2240	2240

Figure 1 - Time series of market one-vega P&L and of the S&P 500 index. This figure plots the time series of the market volatility factor (market one-vega P&L) and of the S&P 500 index. The market one-vega P&L is built from returns of total-delta hedged options written on the S&P 500 index.

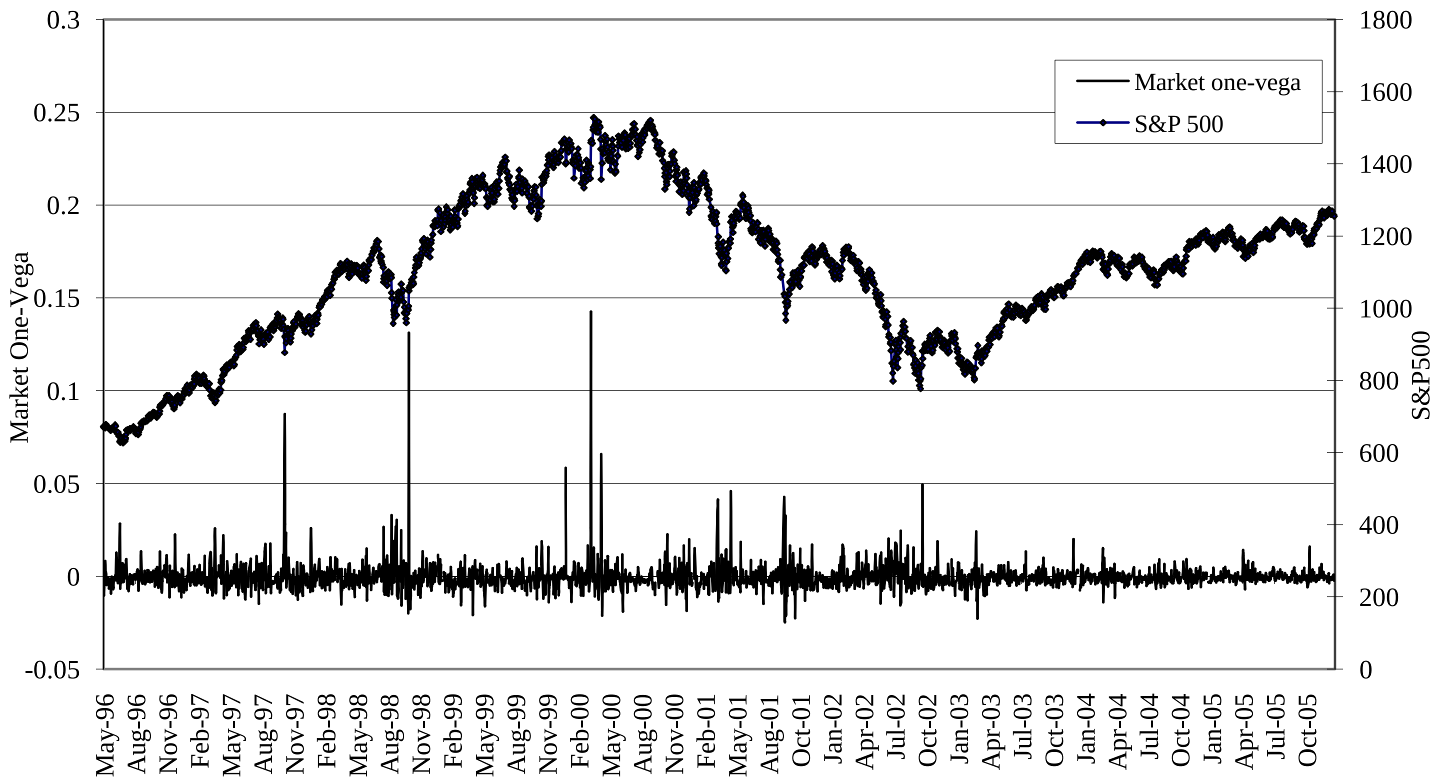

Figure 2 – Option price-error distribution. This figure plots the option price-error distribution used in the simulations. The bid-ask spread of an option in the simulation is set equal to $e^{M+S\times\mu_i}$ where M and S are functions of the time-to-maturity and moneyness of the option, while μ_i is a normally distributed random number for the underlying stock. Option prices in the simulations with measurement errors are equal to the model price plus a random number generated from the plotted probability density.

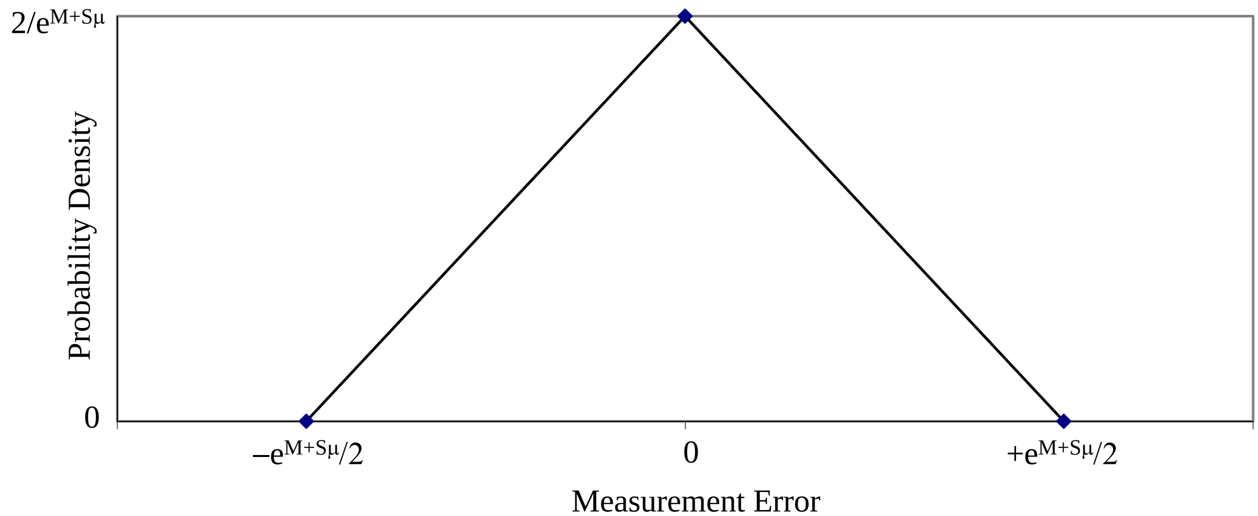

Figure 3 – The relation between volatility betas and average total delta-hedged returns. This figure plots the average returns and volatility betas of relative price-weighted portfolios that are sorted on the basis of maturity, moneyness, and pre-ranking volatility betas. Puts and calls comprise separate portfolios. The sample period is 1997 to 2006. Options with bid-ask spread larger than 25% of the mid option price are deleted. Each letter represents one portfolio, with C for call portfolios and P for put portfolios. Black letters correspond to out-of-the-money options, while gray letters denote in-the-money or at-the-money options.

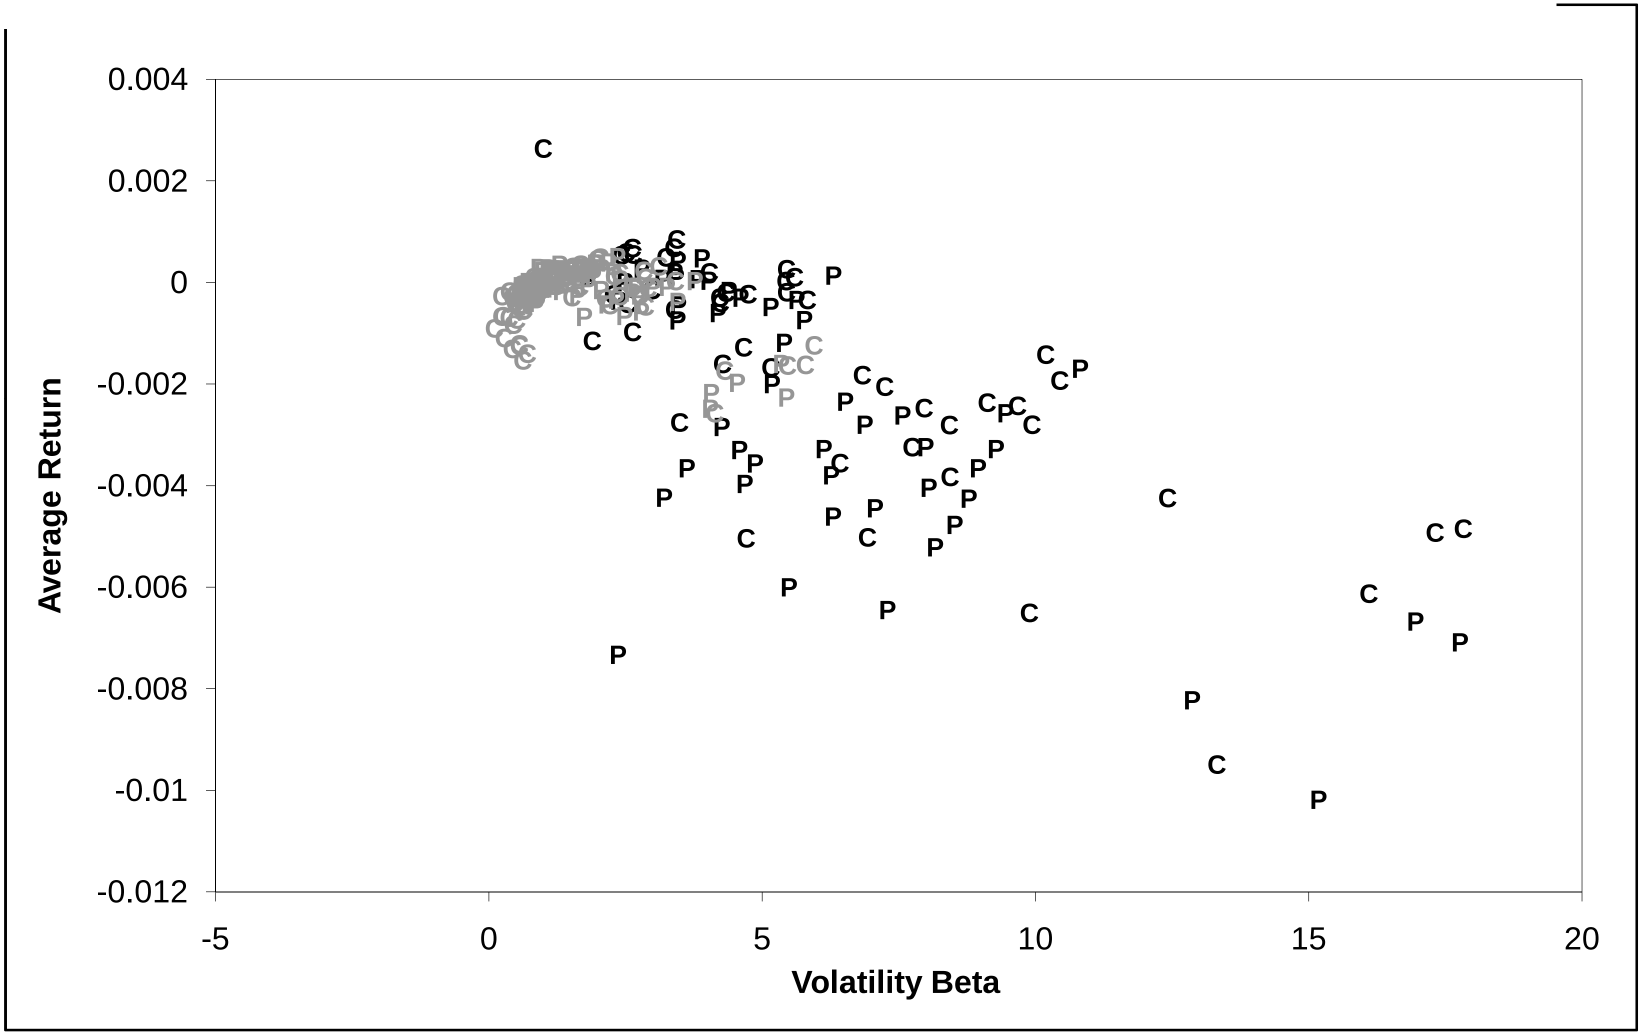

Figure 4 - Fama-MacBeth estimates of realized volatility risk premia. This figure displays the estimated realized volatility risk premia from cross-sectional regressions on beta with the respect to market volatility ($_{Vol}$), the beta with respect to the stock market return ($_{M}$), and bias control variables described in Table 8. Results use OLS regression on 420 equal-weighted portfolios formed on the basis of moneyness, maturity, and pre-ranking volatility beta. Calls and puts comprise separate portfolios. The sample period is 1997 to 2006. Options with bid-ask spread larger than 25% of the mid option price are deleted. The dependent variables are total delta-hedged returns.