

Delta-Hedged Gains and the Negative Market Volatility Risk Premium

Gurdip Bakshi

University of Maryland

Nikunj Kapadia

University of Massachusetts–Amherst

We investigate whether the volatility risk premium is negative by examining the statistical properties of delta-hedged option portfolios (buy the option and hedge with stock). Within a stochastic volatility framework, we demonstrate a correspondence between the sign and magnitude of the volatility risk premium and the mean delta-hedged portfolio returns. Using a sample of S&P 500 index options, we provide empirical tests that have the following general results. First, the delta-hedged strategy underperforms zero. Second, the documented underperformance is less for options away from the money. Third, the underperformance is greater at times of higher volatility. Fourth, the volatility risk premium significantly affects delta-hedged gains, even after accounting for jump fears. Our evidence is supportive of a negative market volatility risk premium.

The notion that volatility of equity returns is stochastic has a firm footing in financial economics. However, a less than understood phenomenon is whether volatility risk is compensated, and whether this compensation is higher or lower than the risk-free rate. Is the risk from changes in market volatility positively correlated with the economy-wide pricing kernel process? If so, how does it affect the equity and option markets? Evidence that market volatility risk premium may be nonzero can be motivated by three empirical findings:

Purchased options are hedges against significant market declines. This is because increased realized volatility coincides with downward market moves [French, Schwert, and Stambaugh (1987) and Glosten, Jagannathan,

We thank Doron Avramov, Mark Broadie, Charles Cao, Peter Carr, Bent Christensen, Sanjiv Das, Stephen Figlewski, Paul Glasserman, Steve Heston, Bob Jarrow, Christopher Jones, Nengjui Ju, Hossein Kazemi, Leonid Kogan, Dilip Madan, George Martin, Vasant Naik, Maureen O'Hara, Jun Pan, Jay Patel, Allen Poteshman, N. R. Prabhala, Lemma Senbet, Rangarajan Sundaram, Bob Whaley, and Gregory Willette for helpful comments and discussions. Parts of the article build on N. Kapadia's thesis written at New York University. Earlier versions of the article were presented at Boston University, University of Massachusetts, and Virginia Polytechnic Institute. Participants at the 1998 WFA (Monterey) meetings, 2001 AFA (New Orleans) meetings, and 11th Annual Derivative Securities Conference provided many useful suggestions. Nick Bollen (AFA discussant) and Jeff Fleming (WFA discussant) provided extremely constructive comments. The reports of Bernard Dumas (the editor) and two referees have substantially improved this article. We are especially grateful for the input of Andrea Buraschi. Bent Christensen graciously shared his option dataset. Kristaps Līcis provided excellent research assistance. N. Kapadia acknowledges financial support from the Center of International Securities and Derivatives Markets. The 1998 version of this article was circulated under the title "Do Equity Options Price Volatility Risk?" Address correspondence to: Nikunj Kapadia, Department of Finance and Operations Management, Isenberg School of Business, University of Massachusetts, Amherst, MA 01003, or e-mail: nkapadia@som.umass.edu.

and Runkle (1993)]. One economic interpretation is that buyers of market volatility are willing to pay a premium for downside protection. The hedging motive is indicative of a negative volatility risk premium.

At-the-money Black–Scholes implied volatilities are systematically and consistently higher than realized volatilities [Jackwerth and Rubinstein (1996)]. A potential explanation for this puzzling empirical regularity is that the volatility risk premium is negative. *Ceteris paribus*, a negative volatility risk premium increases the risk-neutral drift of the volatility process and thus raises equity option prices.

Equity index options are nonredundant securities [Bakshi, Cao, and Chen (2000) and Buraschi and Jackwerth (2001)]. Index option models omitting the economic impact of a market volatility risk premium may be inconsistent with observed option pricing dynamics.

This article investigates, both theoretically and empirically, whether the volatility risk premium is negative in index option markets. This is done without imposing any prior structure on the pricing kernel and without parameterizing the evolution of the volatility process. The setup is a portfolio of a long call position, hedged by a short position in the stock, such that the net investment earns the risk-free interest rate. The central idea underlying our analysis is that if option prices incorporate a nonzero volatility risk premium, then we can infer its existence from the returns of an option portfolio that has dynamically hedged all risks *except* volatility risk.

If volatility is constant, or the price process follows a one-dimensional Markov diffusion, then our theoretical analysis implies that the net gain, henceforth “delta-hedged gains,” on the delta-hedged portfolio is precisely zero. A similar conclusion obtains when volatility is stochastic, but volatility risk is unpriced. In this particular case, we show that the distribution of the delta-hedged gains has an expected value of zero. However, if volatility risk is priced, then the sign and magnitude of average delta-hedged gains are determined by the volatility risk premium.

Our theoretical characterizations point to empirical implications that can be tested using discrete delta-hedged gains. First, in the time series, the *at-the-money* delta-hedged gains must be related to the volatility risk premium. Second, cross-sectional variations in delta-hedged gains (in the strike dimension) are restricted by the sensitivity of the option to volatility (i.e., the vega). Our framework allows us to differentiate between the volatility risk premium and the jump-fear underpinnings of delta-hedged gains.

We test model implications using options written on the S&P 500 index. Our empirical specifications are supportive of the following general results:

- The delta-neutral, positive vega strategy that buys calls and hedges with the underlying stock significantly underperforms zero. On average, over all strike and maturity combinations, the strategy loses about 0.05% of

the market index, and about 0.13% for at-the-money calls. This underperformance is also economically large; for at-the-money options, this amounts to 8% of the option value.

- Cross-sectional regression results indicate that delta-hedged gains are negatively related to vega, after controlling for volatility and option maturity. Consistent with our predictions, the hedged gains are maximized for at-the-money options. Controlling for moneyness, the underperformance is greater when the hedging horizon is extended. These results are robust across time, and to the inclusion of put options.
- At times of higher volatility, this underperformance is even more negative. The losses on long call positions persist throughout the sample period and cannot be reconciled by a downward trending market volatility.

Both the cross-sectional and time-series tests provide evidence that support the hypothesis of a nonzero volatility risk premium.¹ In particular, the results suggest that option prices reflect a *negative* market volatility risk premium. To confirm the hedging rationale underlying a negative volatility risk premium, we empirically estimate the option vega, and verify that it is strictly positive. Moreover, we show that options become more expensive (as measured by implied volatility) after extreme market declines.

Negative delta-hedged gains could be consistent with jump risk. This hypothesis is explored in three ways. First, we examine whether measures of asymmetry and peakedness of the risk-neutral distribution (jump-fear surrogates) are linked to losses on delta-hedged strategies. We find that, while risk-neutral skewness helps explain some portion of delta-hedged gains for short-dated options, the volatility risk premium is the predominant explanatory factor for delta-hedged gains at all maturities. Second, we investigate delta-hedged gains in a precrash sample, and again find that the average excess return on delta-hedged portfolio is negative for both calls and puts. Finally, we observe that delta-hedged gains for options bought immediately after a tail event are not substantially different for negative versus positive returns. In summary, the jump-fear explanation, although plausible, cannot be the sole economic justification for systematic losses incurred on delta-hedged

¹ Coval and Shumway (2001) and Buraschi and Jackwerth (2001) provide additional evidence on the possible existence of a nonzero volatility risk premium. For instance, the statistical examination in Buraschi and Jackwerth supports stochastic models with multiple priced factors. Our article differs from existing treatments in several respects. First, we provide an analytical characterization that links the distribution of the gains on a delta-hedged option portfolio to the underlying risk sources. Specifically we show a correspondence between the sign of the mean delta-hedged gains and the sign of the volatility risk premium. Economically the magnitude of the market volatility risk premium is connected to the value of the option as a hedge. Second, our modeling framework provides an explicit set of hypotheses for testing whether the market volatility risk premium is negative. Instrumental to this thrust is whether priced volatility risks or jumps are the primary source of the underperformance of delta-hedged portfolios. For related innovations, we refer the reader to Jackwerth and Rubinstein (1996), Bakshi, Cao and Chen (1997), Dumas, Fleming, and Whaley (1998), Potesman (1998), Benzoni (1999), Bates (2000), Chernov and Ghysels (2000), Heston and Nandi (2000), Buraschi and Jackwerth (2001), Anderson, Benzoni, and Lund (2002), Chernov et al. (2002), Duan, Popova, and Ritchken (2002), Jones (2002), Pan (2002), and Eraker, Johannes, and Polson (2003).

portfolios. However, stochastic volatility models with a negative volatility risk premium show promise in reconciling this observation.

The rest of the article is organized as follows. Section 1 formulates our theoretical analysis of delta-hedged gains in a stochastic volatility and jump setting. Section 2 discusses the data and variable definitions. The statistical properties of delta-hedged gains are described in Section 3. Sections 4 and 5, respectively, examine the cross-sectional and time-series implications between delta-hedged gains and the volatility risk premium. Section 6 investigates whether volatility risk premium explains delta-hedged gains in the presence of jump factors. Section 7 concludes. All technical details are in the appendix.

1. Delta-Hedged Gains and the Volatility Risk Premium

This section describes the distribution of the gain on a portfolio of a long position in an option, hedged by a short position in the underlying stock, such that the net investment earns the risk-free interest rate. We call the gain on the hedged portfolio the “delta-hedged gains.” We first develop the relevant theory in Section 1.1, assuming that volatility of equity returns is constant, and then relax this assumption in Section 1.2 by allowing volatility to be stochastic. The theoretical implications are then used in Section 1.3 to motivate empirical tests about the negative volatility risk premium. Our analysis also shows how the presence of jumps can contribute to the underperformance of delta-hedged equity portfolios.

To formalize main ideas, let $C(t, \tau; K)$ represent the price of a European call maturing in τ periods from time t , with strike price K . Denote the corresponding option delta by $\Delta(t, \tau; K)$. Define the delta-hedged gains, $\Pi_{t, t+\tau}$, as the gain or loss on a delta-hedged option position, where the net (cash) investment earns the risk-free rate

$$\Pi_{t, t+\tau} \equiv C_{t+\tau} - C_t - \int_t^{t+\tau} \Delta_u dS_u - \int_t^{t+\tau} r(C_u - \Delta_u S_u) du, \tag{1}$$

where S_t is the time t price of the underlying (nondividend paying) stock and r is the constant risk-free interest rate. In Equation (1), we have used the shorthand notation $C_t \equiv C(t, \tau)$ and $\Delta_t \equiv \Delta(t, \tau) = \partial C_t / \partial S_t$ for compactness. The expected $\Pi_{t, t+\tau}$ can be interpreted as the excess rate of return on the delta-hedged option portfolio. Relating delta-hedged gains to the volatility risk premium is our primary objective throughout.

1.1 Delta-hedged gains under constant volatility

Let the stock price follow a geometric Brownian motion under the physical probability measure (with constant drift, μ , and constant volatility, σ):

$$\frac{dS_t}{S_t} = \mu dt + \sigma dW_t^1. \tag{2}$$

By Itô's lemma, we can write the call price as

$$C_{t+\tau} = C_t + \int_t^{t+\tau} \Delta_u dS_u + \int_t^{t+\tau} \left(\frac{\partial C_u}{\partial u} + \frac{1}{2} \sigma^2 S_u^2 \frac{\partial^2 C_u}{\partial S_u^2} \right) du. \quad (3)$$

Standard assumptions also show that the call option price is a solution to the Black–Scholes valuation equation,

$$\frac{1}{2} \sigma^2 S^2 \frac{\partial^2 C}{\partial S^2} + rS \frac{\partial C}{\partial S} + \frac{\partial C}{\partial t} - rC = 0. \quad (4)$$

Using Equations (3) and (4), it follows that

$$C_{t+\tau} = C_t + \int_t^{t+\tau} \Delta_u dS_u + \int_t^{t+\tau} r(C_u - \Delta_u S_u) du, \quad (5)$$

which is the statement that the call option can be replicated by trading a stock and a bond. Combining Equation (5) with the definition of delta-hedged gains in Equation (1), it is apparent that, with continuous trading, $\Pi_{t,t+\tau} = 0$ over every horizon τ . More generally, it can be verified that $\Pi_{t,t+\tau} = 0$ is a property common to all one-dimensional Markov Itô price processes: $\frac{dS_t}{S_t} = \mu_t[S_t] dt + \sigma_t[S_t] dW_t^1$.

When the hedge is rebalanced discretely, $\Pi_{t,t+\tau}$ will not necessarily be zero. However, Bertsimas, Kogan, and Lo (2000) show that the delta-hedged gains have an asymptotic distribution that is symmetric with zero mean. Consider a portfolio of an option that is hedged discretely N times over the life of the option, where the hedge is rebalanced at each of the dates t_n , $n = 0, 1, \dots, N-1$ (where we define $t_0 = t$, $t_N = t + \tau$) and kept in place over the period $t_{n+1} - t_n = \tau/N$. Define the discrete delta-hedged gains, $\pi_{t,t+\tau}$, as

$$\pi_{t,t+\tau} \equiv C_{t+\tau} - C_t - \sum_{n=0}^{N-1} \Delta_{t_n} (S_{t_{n+1}} - S_{t_n}) - \sum_{n=0}^{N-1} r(C_t - \Delta_{t_n} S_{t_n}) \frac{\tau}{N}, \quad (6)$$

where $\Delta_{t_n} \equiv \partial C_{t_n} / \partial S_{t_n}$. From Bertsimas, Kogan, and Lo (2000), $\sqrt{N} \pi \Rightarrow \frac{1}{\sqrt{2}} \int_t^{t+\tau} \sigma^2 S_u^2 \frac{\partial^2 C_u}{\partial S_u^2} dW_u$, where W_u is a Wiener process, independent of W_u^1 . Thus the asymptotic distribution of the discretely hedged option portfolio has a mean of zero and is symmetric. Simulation results in Bertsimas, Kogan, and Lo as well as those reported in Figlewski (1989) suggest that the distribution of $\pi_{t,t+\tau}$ is centered around zero for a wide range of parameters and for low values of N (about 10).

We show next that if we relax the geometric Brownian motion assumption for the stock price and allow for stochastic volatility outside of the one-dimensional Markov diffusion context, then $\pi_{t,t+\tau}$ is centered around zero unless volatility risk is priced. Therefore this setting allows us to construct a test to examine whether the volatility risk premium is negative.

1.2 Delta-hedged gains under stochastic volatility

Consider a (two-dimensional) price process that allows stock return volatility to be stochastic (under the physical probability measure),

$$\frac{dS_t}{S_t} = \mu_t[S_t, \sigma_t] dt + \sigma_t dW_t^1, \tag{7}$$

$$d\sigma_t = \theta_t[\sigma_t] dt + \eta_t[\sigma_t] dW_t^2, \tag{8}$$

where the correlation between the two Weiner processes, W_t^1 and W_t^2 , is ρ . It may be noted that volatility, σ_t , follows an autonomous stochastic process; the drift coefficient, $\theta_t[\sigma_t]$, and the diffusion coefficient, $\eta_t[\sigma_t]$, are functionally independent of S_t . Therefore by Itô's lemma,

$$C_{t+\tau} = C_t + \int_t^{t+\tau} \frac{\partial C_u}{\partial S_u} dS_u + \int_t^{t+\tau} \frac{\partial C_u}{\partial \sigma_u} d\sigma_u + \int_t^{t+\tau} b_u du, \tag{9}$$

where defining $b_u \equiv \frac{\partial C_u}{\partial u} + \frac{1}{2} \sigma_u^2 S_u^2 \frac{\partial^2 C_u}{\partial S_u^2} + \frac{1}{2} \eta_u^2 \frac{\partial^2 C_u}{\partial \sigma_u^2} + \rho \eta_u \sigma_u S_u \frac{\partial^2 C_u}{\partial S_u \partial \sigma_u}$. The valuation equation that determines the price of the call option is

$$\begin{aligned} & \frac{1}{2} \sigma_t^2 S^2 \frac{\partial^2 C}{\partial S^2} + \frac{1}{2} \eta_t^2 \frac{\partial^2 C}{\partial \sigma^2} + \rho \eta_t \sigma_t S \frac{\partial^2 C}{\partial S \partial \sigma} + rS \frac{\partial C}{\partial S} \\ & + (\theta_t - \lambda_t[\sigma]) \frac{\partial C}{\partial \sigma} + \frac{\partial C}{\partial t} - rC = 0, \end{aligned} \tag{10}$$

where $\lambda_t[\sigma] \equiv -\text{cov}_t(\frac{dm_t}{m_t}, d\sigma_t)$ represents the price of volatility risk, for a pricing kernel process m_t , and $\text{cov}_t(\cdot, \cdot)$ is a conditional covariance operator divided by dt . In general, the volatility risk premium will be related to risk aversion and to the factors driving the pricing kernel process. Rearranging Equation (10), it follows that b_u is also equal to

$$b_u = r \left(C_u - S_u \frac{\partial C_u}{\partial S_u} \right) - (\theta_u[\sigma_u] - \lambda_u[\sigma_u]) \frac{\partial C_u}{\partial \sigma_u}. \tag{11}$$

Substituting for b_u in the stochastic differential equation [Equation (9)], we obtain

$$\begin{aligned} C_{t+\tau} = C_t & + \int_t^{t+\tau} \frac{\partial C_u}{\partial S_u} dS_u + \int_t^{t+\tau} \frac{\partial C_u}{\partial \sigma_u} d\sigma_u \\ & + \int_t^{t+\tau} \left(r \left(C_u - S_u \frac{\partial C_u}{\partial S_u} \right) - (\theta_u[\sigma] - \lambda_u[\sigma]) \frac{\partial C_u}{\partial \sigma_u} \right) du. \end{aligned} \tag{12}$$

This equation can be further simplified by substituting for $d\sigma_u$ in the second integral to give

$$C_{t+\tau} = C_t + \int_t^{t+\tau} \frac{\partial C_u}{\partial S_u} dS_u + \int_t^{t+\tau} r \left(C_u - \frac{\partial C_u}{\partial S_u} S_u \right) du + \int_t^{t+\tau} \lambda_u \frac{\partial C_u}{\partial \sigma_u} du + \int_t^{t+\tau} \eta_u \frac{\partial C_u}{\partial \sigma_u} dW_u^2. \quad (13)$$

We are now ready to prove the following relationships between delta-hedged gains, the volatility risk premium, and the option vega.

Proposition 1. *Let the stock price process follow the dynamics given in Equations (7) and (8). Moreover, suppose the volatility risk premium is of the general form $\lambda_t[\sigma_t]$. Then,*

1. *The delta-hedged gains, $\Pi_{t,t+\tau}$, is given by*

$$\Pi_{t,t+\tau} = \int_t^{t+\tau} \lambda_u[\sigma_u] \frac{\partial C_u}{\partial \sigma_u} du + \int_t^{t+\tau} \eta_u[\sigma_u] \frac{\partial C_u}{\partial \sigma_u} dW_u^2, \quad (14)$$

and from the martingale property of the Itô integral

$$E_t(\Pi_{t,t+\tau}) = \int_t^{t+\tau} E_t \left(\lambda_u[\sigma_u] \frac{\partial C_u}{\partial \sigma_u} \right) du, \quad (15)$$

where $\frac{\partial C_t}{\partial \sigma_t}$ represents the vega of the call option, and $E_t(\cdot)$ is the expectation operator under the physical probability measure.

2. *If volatility risk is not priced in equilibrium, that is, $\lambda_t[\sigma_t] \equiv 0$, then*

$$E_t(\pi_{t,t+\tau}) = O(1/N), \quad (16)$$

where the discrete delta-hedged gain, $\pi_{t,t+\tau}$, is as defined previously in Equation (6).

The proposition states that, with continuous trading, if volatility risk is not priced, the delta-hedged gains are, on average, zero. In practice, the hedge is rebalanced discretely over time, and this may bias the average $\pi_{t,t+\tau}$ away from zero. However, in Equation (16), we show that this bias is small. That is, even if we allow for discrete trading, for both the Black–Scholes model and the stochastic volatility model, the mean delta-hedged gain is zero, up to terms of $O(1/N)$.² If volatility risk is priced, Equation (15) shows that $E_t(\Pi_{t,t+\tau})$ is determined by the price of volatility risk, λ_t , and the vega of

² The distribution of the delta-hedged gains can be described in terms of single and multiple Itô integrals. It is difficult to represent multiple Itô integrals in increments of their component Wiener processes [Milstein (1995)]. Therefore, unlike the Black–Scholes case, the asymptotic distribution of $\sqrt{N}\pi_{t,t+\tau}$ cannot be described succinctly.

the option, $\partial C_t / \partial \sigma_t$. The statistical tests in Buraschi and Jackwerth (2001) support a nonzero volatility risk premium.

There are two specific testable implications that follow from Equation (15). First, as the vega is positive, a negative (positive) λ_t implies that $E_t(\Pi_{t,t+\tau})$ will be negative (positive). In particular, a negative volatility risk premium is consistent with the notion that market volatility often rises when the market return drops. To see this, consider a Lucas–Rubinstein investor that is long the market portfolio and has a coefficient of relative risk aversion γ . Under this particular assumption, the pricing kernel $m_t = S_t^{-\gamma}$. An application of Itô’s lemma yields $\lambda_t[\sigma_t] = \gamma \text{cov}_t(dS_t/S_t, d\sigma_t)$, so that a negative correlation between the stock return and the volatility process implies a negative λ_t . In our modeling paradigm, there is a one-to-one correspondence between the sign of λ_t and the sign of mean delta-hedged gains.

Economically, purchased options are hedges against market declines because increased realized volatility tends to occur when market falls significantly. Consequently, in the stochastic volatility setting, the underperformance of the delta-hedged portfolio is tantamount to the existence of a negative volatility risk premium. Our framework allows one to determine the sign of the volatility risk premium without imposing any strong restrictions on the pricing kernel process, and it also does not rely on the identification or the estimation of the volatility process. The quantitative strategy of Equation (6) is relatively easy to implement in option markets.

Second, as the option vega is largest for near-the-money options, the absolute value of $E_t(\Pi_{t,t+\tau})$ is also largest for near-the-money options. If, in addition, the volatility risk premium is negative, as we have hypothesized, the underperformance of the delta-hedged portfolio should decrease for strikes away from at-the-money. We may note that, as option vega’s are negligible, especially for deep in-the-money options, these options have little to say about the nature of the volatility risk premium. Although our focus has been on calls, all the results also apply to puts.

1.3 Testable predictions

Before we can derive precise empirical implications from Equation (15), we need to simplify the right-hand side, $E_t \int_t^{t+\tau} \lambda_u \partial C_u / \partial \sigma_u du$, in terms of the contemporaneous stock price and the level of volatility. To streamline discussion, define $g(S_t, \sigma_t) \equiv \lambda_t[\sigma_t] \frac{\partial C_t}{\partial \sigma_t}$, and consider the Itô–Taylor expansion of $\int_t^{t+\tau} g(S_u, \sigma_u) du$ [Milstein (1995)]:

$$\begin{aligned} & \int_t^{t+\tau} g(S_u, \sigma_u) du \\ &= g(S_t, \sigma_t) \int_t^{t+\tau} du + \int_t^{t+\tau} \int_t^u \mathcal{L}[g(S_u, \sigma_u)] du' du \\ & \quad + \int_t^{t+\tau} \left[\int_t^u \Gamma_1[g(S_u, \sigma_u)] dW^1 + \int_t^u \Gamma_2[g(S_u, \sigma_u)] dW^2 \right] du, \end{aligned}$$

where the (infinitesimal) operators are defined as $\mathcal{L}[\cdot] = \frac{\partial}{\partial t}[\cdot] + \mu_t S_t \frac{\partial}{\partial S_t}[\cdot] + \theta_t \frac{\partial}{\partial \sigma_t}[\cdot] + \frac{1}{2} \sigma_t^2 S_t^2 \frac{\partial^2}{\partial S_t^2}[\cdot] + \frac{1}{2} \eta_t^2 \frac{\partial^2}{\partial \sigma_t^2}[\cdot] + \sigma_t S_t \eta_t \frac{\partial^2}{\partial S_t \partial \sigma_t}[\cdot]$, $\Gamma_1[\cdot] = \sigma_t S_t \frac{\partial}{\partial S_t}[\cdot]$, and $\Gamma_2[\cdot] = \eta_t \frac{\partial}{\partial \sigma_t}[\cdot]$. Using the martingale property of the Itô integral and Equation (15), we therefore have

$$E_t(\Pi_{t,t+\tau}) = g(S_t, \sigma_t) \tau + E_t \int_t^{t+\tau} \int_t^u \mathcal{L}[g(S_{u'}, \sigma_{u'})] du' du, \quad (17)$$

$$= \sum_{n=0}^{\infty} \frac{\tau^{1+n}}{(1+n)!} \mathcal{L}^n \left[\lambda_t \frac{\partial C_t}{\partial \sigma_t} \right], \quad (18)$$

by a recursive application of the Itô–Taylor expansion. Observe that $E_t(\Pi_{t,t+\tau})$ is abstractly related to the current stock price and volatility, and the parameters of the option price, especially the maturity and moneyness. To develop testable empirical specifications, we now exploit certain option properties and derive the functional form of each term in Equation (18). In the discussions that follow, we assume that all parameters of the option price are held fixed.

For a broad class of option models, the call price is homogeneous of degree one in the stock price and the strike price [Merton (1973)]. So for a fixed moneyness, the call price scales with the price of the underlying asset S_t . In this case the option vega, $\partial C_t / \partial \sigma_t$, also scales with S_t . We may therefore separate $g(S_t, \sigma_t) = \alpha_t(\sigma_t, \tau; y) S_t$, for option moneyness y , and $\alpha_t(\cdot)$ independent of S_t . If $\partial C_t / \partial \sigma_t$, and therefore $g(S_t, \sigma_t)$ scales with S_t and volatility risk is priced, then we assert that $\Pi_{t,t+\tau}$ also scales with S_t , where the scaling factor is a function of σ_t (and other parameters of the option contract). To prove this we make use of Equation (18), the standard assumption that the stock price, S_t , follows a proportional stochastic process and the following property of \mathcal{L} as it operates on a function g :

Lemma 1. *Consider $g(S_t, \sigma_t) = \alpha_t(\sigma_t, \tau; y) S_t^\phi$, for any $\phi \in \mathcal{R}$, where α_t is at most a function of maturity and volatility. If S_t obeys a proportional stochastic process, then $\mathcal{L}^n[g(S, \sigma)]$ is also proportional to S_t^ϕ , for all $n \in \{1, 2, 3, \dots\}$.*

The lemma, which is proved in the appendix, shows that if $g(S_t, \sigma_t)$ scales with S_t , so does $\mathcal{L}^n[g]$. It follows from Equation (18) that $\Pi_{t,t+\tau}$ is also proportional to S_t . Thus we can represent $E_t(\Pi_{t,t+\tau})$ as

$$E_t(\Pi_{t,t+\tau}) = S_t \times f_t[\sigma_t, \tau; y], \quad (19)$$

for some $f_t[\cdot]$ that is determined by the functional dependence of λ_t and $\partial C_t / \partial \sigma_t$ on σ_t , and the parameters of the option price, in particular, the option moneyness and maturity. That $E_t(\Pi_{t,t+\tau}) / S_t$ varies in the time series with physical volatility, σ_t , and in the cross section (for a fixed σ_t) with the option moneyness, y , forms the basis of the empirical tests.

To derive the cross-sectional test, we keep σ_t as fixed, and write $E_t(\Pi_{t,t+\tau}^i)/S_t = f_t[\tau; y_{i,t}]$, for moneyness corresponding to strike price K_i . It is important to keep σ_t as fixed, as the option price is nonlinear in σ_t for away-from-the-money strikes. In the absence of any information regarding the form of the nonlinearity, it is difficult to specify a model and the corresponding econometric test that allows both σ_t and $y_{i,t}$ to vary simultaneously. Given a suitable model for $f_t[\tau; y]$, we can then test the relation between $E_t(\Pi_{t,t+\tau}^i)/S_t$ and $y_{i,t}$. Because the vega of the option and thus the absolute value of $E_t(\Pi_{t,t+\tau}^i)$ is, ceteris paribus, maximized for at-the-money options, and decreases for strikes away from at-the-money, it follows that $f_t[y]$ must also be of such a functional form (controlling for volatility). We can reject the hypothesis of a nonzero volatility risk premium if we do not find this hypothesized relation between $E_t(\Pi_{t,t+\tau}^i)/S_t$ and $y_{i,t}$. Thus the cross section of delta-hedged gains contains information about the volatility risk premium.

Next, to develop the time-series relation between $E(\Pi_{t,t+\tau})/S_t$ and σ_t , consider Equation (19) applied to *at-the-money options*. It has been noted elsewhere [Stein (1989)] that the short-term at-the-money call is almost linear in volatility. If C_t is linear in σ_t , $\partial C_t/\partial \sigma_t$ will be independent of σ_t , and the functional dependence of $\Pi_{t,t+\tau}$ on σ_t will be determined only by λ_t and the underlying stochastic volatility process. Given the functional form of $\lambda_t[\sigma_t]$ and the underlying volatility process, we can infer the functional form of $\Pi_{t,t+\tau}$. For at-the-money options, we may specialize $f_t[\sigma_t, \tau; y] = \hat{f}_t[y; \tau]\hat{f}_t[\sigma_t]$. To make this point precise, we develop the functional form of at-the-money $\Pi_{t,t+\tau}$ for the Heston (1993) model. In his model, the volatility risk premium is linear in volatility [see also the set of assumptions in Bates (2000), Pan (2000), and Eraker, Johannes, and Polson (2003)].

Proposition 2. Consider the special case of the stock price process in Equations (7) and (8), where $\theta[\sigma_t] = -\kappa\sigma_t$ and $\eta[\sigma_t] = v$. Specifically,

$$\frac{dS_t}{S_t} = \mu_t[S_t, \sigma_t]dt + \sigma_t dW_t^1, \tag{20}$$

$$d\sigma_t = -\kappa\sigma_t dt + v dW_t^2, \tag{21}$$

and the volatility risk premium is linear in volatility, as in $\lambda_t[\sigma_t] = \lambda\sigma_t$. Let the call option vega be proportional to S_t and independent of σ_t , as in $\partial C_t/\partial \sigma_t = \beta_t(\tau; y)S_t$. Then the delta-hedged gains for near-the-money options must be

$$E_t(\Pi_{t,t+\tau}) = \lambda\varphi_t(\tau)S_t\sigma_t, \tag{22}$$

where $\varphi_t(\tau) > 0$ is defined in the appendix. At-the-money delta-hedged gains are negative only if $\lambda < 0$.

Specifically for at-the-money options, Proposition 2 shows that if λ_t is proportional to σ_t , so is the scaled delta-hedged gains, $E_t(\Pi_{t,t+\tau})/S_t$. Although not done here, it is straightforward to extend the analysis to other models, in which case, more generally, $E_t(\Pi_{t,t+\tau})/S_t$ may be a polynomial in σ_t [i.e., Hull and White (1987)]. We can thus construct a time-series test relating the scaled at-the-money delta-hedged gains to physical return volatility (or equivalently the volatility risk premium). We can reject the hypothesis of a zero-volatility risk premium if we find a relation between at-the-money $E_t(\Pi_{t,t+\tau})/S_t$ and any function of physical volatility.

In summary, our theoretical results indicate that the bias in $\Pi_{t,t+\tau}$ from discrete hedging is small relative to the impact of a volatility risk premium (as suggested by Proposition 1). Moreover, the mean at-the-money delta-hedged gain (normalized by the stock price) is approximately linear in the level of physical volatility. We verified both these results via simulations. More exactly, the delta-hedged strategy typically underperforms (overperforms) zero with negative (positive) volatility risk premium. In addition, the negative bias is related to the change that occurs because a negative volatility risk premium increases the option price. In large part, the level of underperformance is greater with higher volatility. The details are provided in appendix B.

Before we operationalize and implement the cross-sectional and time-series tests using options data, one question remains unresolved: How is the performance of delta-hedged strategies affected by jumps? To address this question, we appeal to a jump-diffusion model for the equity price [Bates (2000), Merton (1976), and Pan (2000)]. Consider

$$\frac{dS_t}{S_t} = \mu_t[S_t, \sigma_t] dt + \sigma_t dW_t^1 + (e^x - 1) dq_t - \mu_J \Lambda_J \sigma_t dt, \quad (23)$$

where the volatility dynamics are as displayed in Equation (8). This framework allows for both stochastic volatility as well as random jumps to affect delta-hedged gains. The setup is briefly as follows. First, in Equation (23), the variable q_t represents a Poisson jump counter with volatility-dependent intensity $\Lambda_J \sigma_t$. Denote the physical density of the jump size, x , by $q[x]$. Second, we posit that x and q_t are orthogonal to each other and to all sources of uncertainty. In addition, if we assume that the mean of $e^x - 1$ is μ_J , the compensator is $\mu_J \Lambda_J \sigma_t dt$, which is the final term in Equation (23). Lastly, to isolate the impact of jump size and jump intensity on delta-hedged gains, for now we assume that only jump size is priced. The jump risk premium will therefore introduce a wedge between the physical density, $q[x]$, and the risk-neutral density, $q^*[x]$. Specifically, assume that the risk-neutral mean of $e^x - 1$ is μ_J^* .

In the stochastic environment of Equation (23), the delta-hedged gains are equal to (see the appendix):

$$\begin{aligned}
 E_t(\Pi_{t,t+\tau}) = & \int_t^{t+\tau} E_t \left(\lambda_u [\sigma_u] \frac{\partial C_u}{\partial \sigma_u} \right) du + \mu_j^* \Lambda_j \int_t^{t+\tau} E_t \left(\frac{\partial C_u}{\partial S_u} \sigma_u S_u \right) du \\
 & - \Lambda_j \int_t^{t+\tau} \sigma_u du \left\{ \int_{-\infty}^{\infty} C_u(S_u e^x) q^*[x] dx \right. \\
 & \quad \left. - \int_{-\infty}^{\infty} C_u(S_u e^x) q[x] dx \right\}. \tag{24}
 \end{aligned}$$

The first term is a consequence of the volatility risk premium and the other two terms are a consequence of jumps. When $\lambda_u = 0$, Equation (24) imparts the intuition that delta-hedged gains are negative provided the mean jump size is negative (i.e., $\mu_j^* < 0$), and there are occasional price discontinuities (i.e., $\Lambda_j > 0$). In theory, the *fatter left tails* of the equity price distribution can lead to the underperformance of delta-hedged portfolios (the sign of return skewness is determined by the sign of the mean jump size and Λ_j controls excess kurtosis). Equation (24) suggests the effect of jumps on delta-hedged gains is most pronounced for in-the-money options. In our extended framework, the bias in delta-hedged gains is partly due to price volatility risk and partly due to jump exposures.

Observe that the final double integral term in Equation (24) is typically negative. This is because the option price evaluated at the risk-neutral density of the jump size is generally higher than under the physical density. Moreover, when the jump risk premium is volatility dependent, as is the case here, the component of delta-hedged gains due to jump risk is related to variations in volatility. In particular, the higher the physical volatility, the more negative are the total delta-hedged gains. Now, if one additionally assumes that jump intensity is priced (Λ_j gets altered to Λ_j^*), the expression for $E_t(\Pi_{t,t+\tau})$ must be modified. Specifically the last two terms must be replaced by $\mu_j^* \Lambda_j^* \int_t^{t+\tau} E_t \left(\frac{\partial C_u}{\partial S_u} \sigma_u S_u \right) du - \Lambda_j^* \int_t^{t+\tau} \sigma_u \int_{-\infty}^{\infty} \{ C_u(S_u e^x) - C_u(S_u) \} q^*[x] dx du + \Lambda_j \int_t^{t+\tau} \sigma_u \int_{-\infty}^{\infty} \{ C_u(S_u e^x) - C_u(S_u) \} q[x] dx du$. This analysis suggests that both forms of jump risk will lead to the underperformance of delta-hedged portfolios. As we will see, Equation (24) provides the impetus for empirically differentiating between the negative volatility risk premium and the jump fear explanations for negative delta-hedged portfolio returns.

2. Description of Option Data and Variable Definitions

All empirical tests employ daily observations on S&P 500 index options. The option prices consist of time-stamped calls and puts, and correspond to the last bid-ask quote reported before 3:00 PM CST. Rubinstein (1994)

and Jackwerth and Rubinstein (1996) have suggested that the precrash and postcrash index distributions differ considerably. The initial sample date was accordingly chosen to begin from January 1, 1988, to avoid mixing precrash and postcrash options [see also Christensen and Prabhala (1998)]. Our option sample ends on December 30, 1995.

The option universe is constructed in the following way. First, the option data is screened to eliminate option prices that violated arbitrage bounds. Specifically, we exclude call options whose price is outside of the range $(Se^{-z\tau} - e^{-r\tau}K, Se^{-z\tau})$ for dividend yield z . Second, to minimize the impact of recording errors, we discard all options that have Black–Scholes implied volatilities exceeding 100%, or less than 1%. Third, we deleted options with maturity less than 14 days.

In addition, all options with maturities longer than 60 days are eliminated. Our present focus on short-term options allows us to reduce the impact of stochastic interest rates. Finally, deep in-the-money option prices can be unreliable due to the lack of trading volume. As in Jackwerth and Rubinstein (1996) and Buraschi and Jackwerth (2001), deep away-from-the-money options are omitted. Our option sampling procedure results in 36,237 calls and 35,030 puts.

We require a series for dividends, interest rates, and index return volatility. First, following a prevalent practice, we assume that daily dividends are known over the life of the option contract. That is, we take the actual dividends from the S&P 500 bulletin and subtract the present discounted value of dividends from the contemporaneous stock price. The adjusted stock price is employed in our empirical tests throughout.

The interest rates are computed using the procedure outlined in Jackwerth and Rubinstein (1996) and Buraschi and Jackwerth (2001). Each day we infer the interest rate using the put-call parity. Specifically we use strike price and maturity-matched pairs of puts and calls, quoted within a 1-minute interval. The borrowing rate is computed as $r_b = -(1/\tau) \log((Se^{-z\tau} - C^a + P^b)/K)$. Likewise, the lending rate is $r_l = -(1/\tau) \log((Se^{-z\tau} - C^b + P^a)/K)$ for each pair of bid-ask call and put quotes, (C^b, C^a) and (P^b, P^a) . The daily rate used in the tests is the midpoint of r_b and r_l , averaged across all strikes of a specific maturity.

For robustness, we adopt two measures of return volatility. One, we estimate a GARCH(1,1) model using daily S&P 500 returns over the entire period:

$$R_{t-1,t} = \bar{R} + \epsilon_t, \quad (25)$$

$$\sigma_t^2 = a_0 + a_1 \epsilon_{t-1}^2 + a_2 \sigma_{t-1}^2, \quad \text{and}, \quad (26)$$

$$\epsilon_t = \sigma_t \nu_t, \quad \nu_t \sim \text{i.i.d. } \mathcal{N}(0, 1), \quad (27)$$

where the τ -period return is defined as $R_{t,t+\tau} \equiv \log(S_{t+\tau}/S_t)$ and σ_t is the conditional volatility. Relying on the GARCH model estimates, the τ -period GARCH volatility estimate is

$$\text{VOL}_t^g \equiv \sqrt{\frac{252}{\tau} \sum_{n=t-\tau}^t \hat{\sigma}_n^2}, \tag{28}$$

where $\hat{\sigma}_n$ is the fitted value obtained from the GARCH estimation. We experimented with other GARCH specifications and obtained similar volatility estimates. The GARCH volatility measure also allows us to construct a daily volatility series for estimating the hedge ratio in Equation (6).

The other volatility measure is the estimate of the sample standard deviation, as in

$$\text{VOL}_t^h = \sqrt{\frac{252}{\tau} \sum_{n=t-\tau}^t (R_{n-1,n} - \bar{R})^2}, \tag{29}$$

where \bar{R} is now the average daily return. This rolling estimation procedure produces volatility estimates, with estimation error serially uncorrelated through time for nonoverlapping periods.

To construct an empirical test design that limits overlapping observations, we will sometimes appeal to a sample of options with constant maturity (e.g., 30 days and 44 days). Over our sample period, the S&P 500 index options have continual option quotes available only for the two near months. Thus, to build as large a series as possible and yet limit overlap, we employ options of maturity no more than 60 days.

Define the option moneyness as $y \equiv Se^{(r-z)\tau}/K$. Consequently a call (put) option is classified as out-of-the-money if it has moneyness corresponding to $y < 1$ ($y > 1$). For reasons already discussed, our empirical work is restricted to the $\pm 10\%$ moneyness range. To maintain tractability, much of our analysis centers on calls.

3. Statistical Properties of Delta-Hedged Gains

We compute the discrete delta-hedged gains for each call option in two steps: First, at time t , one call is bought at the closing price, C_t . Second, the call is hedged discretely until expiration, $t + \tau$, with the hedge ratio, Δ_t , recomputed daily at the close of the day price. The total delta-hedged gain for each option up to the maturity date is then calculated as

$$\pi_{t,t+\tau} = \max(S_{t+\tau} - K, 0) - C_t - \sum_{n=0}^{N-1} \Delta_{t_n} (S_{t_{n+1}} - S_{t_n}) - \sum_{n=0}^{N-1} r_n (C_t - \Delta_{t_n} S_{t_n}) \frac{\tau}{N},$$

where $t_0 = t$, $t_N = t + \tau$ is the maturity date, and Δ_{t_n} is the hedge ratio at t_n . In our implementation procedure, the interest rate is updated on a daily basis.

For tractability, Δ_{t_n} is computed as the Black–Scholes hedge ratio, $\Delta_{t_n} = \mathcal{N}[d_1(S_{t_n}, t_n)]$, where $\mathcal{N}[\cdot]$ is the cumulative normal distribution and

$$d_1 \equiv \frac{1}{\sigma_{t, t+\tau} \sqrt{\tau_n}} \log(y_n) + \frac{1}{2} \sigma_{t, t+\tau} \sqrt{\tau_n}. \quad (30)$$

All our delta-hedged calculations allow for time-varying volatility, as reflected by the use of GARCH volatility in Equation (30). Although the Black–Scholes hedge ratio is a reasonable estimate of the true hedge ratio when volatility is not correlated with the stock return process, it will be biased otherwise. In a later section we will examine the impact of a misspecified delta.

Panel A of Table 1 provides descriptive statistics for delta-hedged gains grouped over maturity and moneyness combinations. Specifically, we report the averages for (i) dollar delta-hedged gains $\pi_{t, t+\tau}$, (ii) delta-hedged gains scaled by the index level $\pi_{t, t+\tau}/S_t$ (in percent), and (iii) delta-hedged gains scaled by the call price $\pi_{t, t+\tau}/C_t$ (in percent). For at-the-money calls, and for each maturity, the delta-hedging strategy loses money. On average, over all moneyness and maturities, the strategy loses about 0.05% of the index level, and for at-the-money calls (i.e., $y \in [-2.5\%, 2.5\%]$), the strategy loses about 0.10%. Moreover, the mean $\pi_{t, t+\tau}/C_t$ over the full eight-year sample is -12.18% . It may be noted that the reported standard errors, computed as the sample standard deviation divided by the square root of the number of options, are relatively small. The delta-hedged gains are statistically significant in all moneyness and maturity categories.

The average loss on the delta-hedged strategy of about \$0.43 for at-the-money options also appears high compared with the mean bid-ask spread of \$0.375. This finding implies that the buyer of the call (“long” volatility) is paying the seller of the call (“short” volatility) a premium of about \$0.43 per call. The economic impact of this premium is substantial, given the large volume of S&P 500 contracts traded. The S&P 500 trading volume in 1991 was about 11 million contracts, so that the dollar impact of this premium could be as high as \$500 million. The cumulative impact over the eight-year period is on the order of several billion dollars.

We can make two additional empirical observations that appear broadly consistent with a volatility risk premium. First, the mean delta-hedged gains for away-from-the-money strikes are mostly negative, and less so relative to at-the-money calls. Consider options with moneyness $y \in [-7.50\%, -5\%]$ versus options with moneyness $y \in [-2.50\%, 0\%]$. In the “All” category, we can observe that the dollar delta-hedged gains is \$ -0.28 versus \$ -0.42 . Because the vega for away-from-the-money options is small, the impact of the volatility risk premium should be small. Second, the losses on delta-hedged portfolios generally deepen when the hedging horizon is extended from 14–30 days to 31–60 days. For at-the-money options, the dollar loss

Table 1
Delta-hedged gains for S&P 500 index calls

Panel A: Full sample period

Moneyness	\mathcal{N}	π (in \$)			π/S (in %)			π/C (in %)			$1_{\pi < 0}$
		14-30	31-60	All	14-30	31-60	All	14-30	31-60	All	
$y - 1$											
-10% to -7.5%	1284	-0.18 (0.02)	-0.37 (0.03)	-0.31 (0.02)	-0.06 (0.01)	-0.13 (0.01)	-0.11 (0.01)	-97.37 (3.81)	-56.34 (3.50)	-68.30 (2.80)	87
-7.5% to -5%	3619	-0.19 (0.01)	-0.32 (0.02)	-0.28 (0.01)	-0.06 (0.00)	-0.10 (0.01)	-0.09 (0.00)	-70.71 (3.81)	-38.13 (2.08)	-48.05 (1.87)	78
-5% to -2.5%	5684	-0.21 (0.02)	-0.38 (0.02)	-0.32 (0.01)	-0.06 (0.01)	-0.11 (0.01)	-0.09 (0.00)	-34.43 (2.57)	-11.88 (1.54)	-20.19 (1.36)	72
-2.5% to 0%	5903	-0.26 (0.04)	-0.52 (0.03)	-0.42 (0.02)	-0.06 (0.01)	-0.13 (0.01)	-0.10 (0.01)	-6.98 (1.54)	-7.95 (0.72)	-7.59 (0.73)	68
0% to 2.5%	5752	-0.29 (0.04)	-0.55 (0.03)	-0.45 (0.02)	-0.07 (0.01)	-0.13 (0.01)	-0.11 (0.01)	-2.93 (0.51)	-4.44 (0.31)	-3.88 (0.27)	68
2.5% to 5%	5530	-0.05 (0.03)	-0.26 (0.02)	-0.19 (0.02)	-0.01 (0.01)	-0.07 (0.01)	-0.05 (0.01)	-0.31 (0.19)	-1.34 (0.15)	-0.96 (0.12)	60
5% to 7.5%	4811	0.24 (0.02)	0.18 (0.02)	0.20 (0.02)	0.05 (0.01)	0.03 (0.01)	0.04 (0.00)	0.93 (0.10)	0.62 (0.09)	0.74 (0.07)	41
7.5% to 10%	3647	0.36 (0.02)	0.54 (0.02)	0.47 (0.02)	0.08 (0.01)	0.12 (0.01)	0.11 (0.00)	1.03 (0.07)	1.51 (0.07)	1.33 (0.05)	26

Panel B: Delta-hedged gains across the 88:01-91:12 and 92:01-95:12 subsamples

Moneyness	$y - 1$	Sample	\mathcal{N}	π (in \$)			π/S (in %)			π/C (in %)			$1_{\pi < 0}$
				14-30	31-60	All	14-30	31-60	All	14-30	31-60	All	
-10% to -7.5%		Set 1	989	-0.18	-0.47	-0.36	-0.06	-0.17	-0.13	-97.34	-46.83	-65.01	86
		Set 2	295	-0.10	-0.14	-0.13	-0.02	-0.03	-0.03	-97.82	-77.86	-76.73	92
-7.5% to -5%		Set 1	1930	-0.22	-0.45	-0.36	-0.07	-0.16	-0.12	-68.44	-18.98	-38.60	78
		Set 2	1688	-0.11	-0.20	-0.18	-0.03	-0.05	-0.04	-75.88	-54.62	-58.85	79
-5% to -2.5%		Set 1	2308	-0.29	-0.45	-0.39	-0.09	-0.15	-0.13	-19.96	-7.39	-12.19	70
		Set 2	3375	-0.16	-0.34	-0.28	-0.03	-0.08	-0.06	-44.94	-14.84	-25.66	73
-2.5% to 0%		Set 1	2240	-0.18	-0.26	-0.23	-0.06	-0.09	-0.08	-2.39	-0.93	-1.48	63
		Set 2	3668	-0.31	-0.67	-0.54	-0.07	-0.15	-0.12	-9.84	-12.17	-11.31	70
0 to 2.5%		Set 1	2211	-0.08	-0.28	-0.20	-0.03	-0.09	-0.07	-0.10	-1.46	-0.96	63
		Set 2	3540	-0.42	-0.72	-0.61	-0.09	-0.16	-0.13	-4.69	-6.31	-5.71	71
2.5 to 5%		Set 1	2119	-0.02	-0.19	-0.12	-0.01	-0.06	-0.04	0.02	-0.83	-0.51	59
		Set 2	3410	-0.08	-0.31	-0.22	-0.02	-0.07	-0.05	-0.52	-1.65	-1.24	60
5 to 7.5%		Set 1	1882	0.14	0.02	0.06	0.04	0.00	0.02	0.68	0.21	0.39	50
		Set 2	2928	0.31	0.27	0.29	0.07	0.05	0.06	1.10	0.88	0.96	35
7.5 to 10%		Set 1	1523	0.19	0.30	0.26	0.06	0.09	0.07	0.71	1.12	0.95	38
		Set 2	2123	0.50	0.69	0.62	0.10	0.14	0.13	1.30	1.76	1.60	16

We compute the gain on a portfolio of a long position in a call option, hedged by a short position in the underlying stock, such that the net investment earns the risk-free interest rate. The discretely rebalanced delta-hedged gains, $\pi_{t, t+\tau}$, are computed as

$$\pi_{t, t+\tau} = \max(S_{t+\tau} - K, 0) - C_t - \sum_{n=0}^{N-1} \Delta_{t_n} (S_{t_{n+1}} - S_{t_n}) - \sum_{n=0}^{N-1} r_n (C_t - \Delta_{t_n} S_{t_n}) \frac{\tau}{N},$$

where the interest rate, r_n , and the option delta, Δ_{t_n} , are updated on a daily basis. The option delta is computed as the Black-Scholes hedge ratio evaluated at the GARCH volatility. The rebalancing frequency, τ/N , is set to 1 day. We report (i) the dollar delta-hedged gains ($\pi_{t, t+\tau}$), (ii) the delta-hedged gains normalized by the index level ($\pi_{t, t+\tau}/S_t$), and (iii) the delta-hedged gains normalized by the option price ($\pi_{t, t+\tau}/C_t$). All delta-hedged gains are averaged over their respective moneyness and maturity category. The moneyness of the option is defined as $y \equiv Se^{(r-\sigma^2)\tau}/K$. The standard error, shown in parentheses, is computed as the sample standard deviation divided by the square root of the number of observations. $1_{\pi < 0}$ is the proportion of delta-hedged gains with $\pi < 0$, and \mathcal{N} is the number of options. Results are shown separately for options with maturity 14-30 days and 31-60 days; "All" combines the delta-hedged gains from both maturities. There are 36,237 call option observations on the S&P 500 index. Subsample results are displayed in panel B: Set 1 corresponds to 1988:01-1991:12; Set 2 corresponds to 1992:01-1995:12 (standard errors are small and omitted in panel B).

over the 31–60 days maturity is almost twice the loss in the 14- to 30-day maturity. This empirical finding tallies with the theoretical prediction that delta-hedged gains should become more negative with maturity (because the vega is increasing with maturity). Overall the delta-hedged gains are negative except for deep in-the-money options. That deep in-the-money calls have positive delta-hedged gains is anomalous. We will reconcile this result shortly.

Next, to ensure that the documented results are not driven by extremes, we also examine the relative outcomes of positive and negative delta-hedged gains. The last column displays the $1_{\pi < 0}$ statistic which measures the frequency of negative delta-hedged gains (consolidated over all maturities). For at-the-money (out-of-the money) options, it is assuring that 68% (76%) of the observations have negative gains. Therefore the observation that the mean delta-hedged gains are negative, on average, appears robust. Moreover, the frequency of negative delta-hedged gains rise (fall) monotonically when options go progressively out of the money (in the money). If deep in-the-money calls are excluded, then as much as 72% of the remaining call sample have negative delta-hedged gains.

As seen from Panel B of Table 1, the results are robust across subsamples (the standard errors are small and suppressed). In the rows marked Set 1 and Set 2, we report the mean delta-hedged gains over the 88:01–91:12 and the 92:01–95:12 sample periods, respectively. Clearly the underperformance of the delta-hedged strategy is more pronounced over the second subsample. In yet another exercise, we examined the sensitivity of our conclusions to any unexpected declines in index volatility (the delta-hedged portfolios suffer losses when volatility declines). The delta-hedged gains for at-the-money options are negative in seven of eight years. Therefore the persistent losses on the delta-hedged portfolios cannot be attributable to any secular declines in index volatility. Finally, to verify the results from a different options market, we examined delta-hedged gains using options on the S&P 100 index (the details were reported in an earlier version). Reassuringly, the mean delta-hedged gains are also negative for S&P 100 index options. Our conclusions are robust across sample periods as well as across both index option contracts.

Although the conventional estimates of the cross-sectional standard errors are small in both the full sample and the subsamples, these standard errors may not account for the fact that the theoretical distribution of $\pi_{t,t+\tau}$ depends on option moneyness and maturity. We attack this problem on two fronts. First, we construct representative option time series that are homogeneous with respect to moneyness and maturity. Specifically we take at-the-money call options with a fixed maturity of 30 days, 44 days, and 58 days, and delta-hedged them until maturity. For 30-day calls we get a mean $\pi = -\$0.47$ with a t -statistic of -2.34 . Similarly the mean π for 44-(58-) day options is $-0.53(-0.63)$, with a t -statistic of $-2.90 (-2.80)$. Therefore, inferences

based on a homogeneous time series of delta-hedged gains (and standard t -tests) also reject the null hypothesis of zero mean delta-hedged gains.

Second, the standard deviation of discrete delta-hedged gains in the context of one-dimensional diffusions is known from Bertsimas, Kogan, and Lo (2000). For Black–Scholes (see their Theorems 1 and 3), this standard deviation equals $\frac{K\sigma}{2\sqrt{\pi}} [\int_0^1 (1-u^2)^{-1/2} \exp(\frac{(\mu u - \sigma^2/2 + \log(S_t/K))}{\sigma^2(1+u)}) du]^{1/2}$. Even though analytical, the above expression requires estimates of the expected rate of return, μ , and the volatility, σ . We set $\mu = 11.6\%$ and $\sigma = 11\%$ to match the average annual index return and volatility in our sample. For a given strike K , we compute the standard deviation of $\pi_{t,t+\tau}$ at each date t . Standardizing each $\pi_{t,t+\tau}$ by the corresponding standard deviation results in a variable with unit variance. Adhering to a standard practice, we then compute the t -statistic as the average standardized π multiplied by the square root of the number of observations. The resulting t -statistics are -5.29 for 30-day options and -7.30 (-9.32) for 44-day (58-day) options. That the standard deviation of the distribution of $\pi_{t,t+\tau}$ decreases with maturity when the hedge ratio is updated daily is to be expected [see the simulations in Table 1 of Bertsimas, Kogan, and Lo (2000)]. Reinforcing our earlier results, under Black–Scholes, we can easily reject the hypothesis that $E_t(\pi_{t,t+\tau}) = 0$. We tried other combinations of μ and σ , and obtained similar results. It would be of interest to extend this analysis by theoretically characterizing the distribution of delta-hedged gains under stochastic volatility.

Now return to the result that the delta-hedged gains are typically positive for deep in-the-money options, with moneyness greater than 5%. The relative illiquidity for in-the-money calls may upwardly bias the mean delta-hedged gains. Because there is not much trading activity, the market makers often choose not to update in-the-money call prices in response to small changes in the index level [Bakshi, Cao, and Chen (2000)]. It is possible that illiquidity of in-the-money options contribute to positive delta-hedged portfolio returns. To verify this conjecture, and to understand the sources of this phenomena, in Table 2 we examine delta-hedged gains for out-of-the-money puts, which are equivalent to in-the-money calls. Relative to in-the-money calls, the out-of-the-money puts are more actively traded. Supportive of our conjecture, and in contrast to the empirical results from in-the-money calls, the out-of-the-money put delta-hedged gains are now strongly negative:

- The average delta-hedged gains are $-\$1.03$ and $-\$0.82$ for put options with moneyness $y \in [5\%, 7.5\%)$ and $y \in [7.5\%, 10.0\%)$, respectively. It is evident that the losses on the delta-hedged put portfolios is robust to samples restricted by strikes, maturity, and time periods;
- When deep in-the-money calls (beyond 5%) are combined with deep out-of-the-money puts, the mean dollar delta-hedged gains are $\$ - 0.14$, and of absolute magnitude less than that for all at-the-money calls and puts.

Table 2
Delta-hedged gains for out-of-the-money puts

Moneyness $y - 1$	Sample	N	π (in \$)			π/P (in %)			$1_{\pi < 0}$ %
			14-30	30-60	All	14-30	31-60	All	
0-2.5%	Full	5342	-0.55	-0.77	-0.69	-16.95	-13.01	-14.47	74
	Set 1	2116	-0.16	0.10	0.00	0.06	4.82	3.03	62
	Set 2	3226	-0.79	-1.27	-1.09	-27.10	-23.27	-24.68	81
2.5-5%	Full	5695	-0.80	-1.37	-1.16	-55.40	-41.72	-46.82	89
	Set 1	2011	-0.66	-0.80	-0.75	-23.83	-13.95	-17.63	81
	Set 2	3684	-0.88	-1.68	-1.38	-72.61	-56.89	-62.75	94
5-7.5%	Full	5364	-0.60	-1.30	-1.03	-73.57	-66.82	-69.41	94
	Set 1	2042	-0.66	-1.18	-0.97	-45.06	-39.23	-41.56	88
	Set 2	3322	-0.56	-1.38	-1.07	-92.25	-83.11	-86.53	98
7.5-10%	Full	3815	-0.43	-1.06	-0.82	-91.59	-82.66	-86.14	97
	Set 1	1568	-0.57	-1.12	-0.90	-81.72	-60.90	-69.17	94
	Set 2	2247	-0.33	-1.02	-0.76	-98.71	-97.54	-97.99	99

This table reports the delta-hedged gains for out-of-the-money puts on the S&P 500 index. Put options correspond to moneyness, y , greater than 1. We compute the gain on a portfolio of a long position in a put option, hedged by a short position in the underlying stock, such that the net investment earns the risk-free interest rate. As before, the discretely rebalanced delta-hedged gains, $\pi_{t, t+\tau}$, are computed as

$$\pi_{t, t+\tau} = \max(K - S_{t+\tau}, 0) - P_t - \sum_{n=0}^{N-1} \widehat{\Delta}_n (S_{t_{n+1}} - S_{t_n}) - \sum_{n=0}^{N-1} r_n (P_t - \widehat{\Delta}_n S_{t_n}) \frac{\tau}{N},$$

where $\widehat{\Delta}_n$ is the Black-Scholes put option delta evaluated at GARCH volatility, and r_n is the nominal interest rate. The rebalancing frequency, τ/N , is set to 1 day. Reported are (i) dollar delta-hedged gains ($\pi_{t, t+\tau}$) and (ii) delta-hedged gains normalized by the put price ($\pi_{t, t+\tau}/P_t$). All delta-hedged gains are averaged over their respective moneyness and maturity categories. $1_{\pi < 0}$ is the proportion of delta-hedged gains with $\pi < 0$. N represents the number of put options. There are 20,216 out-of-the-money puts. Set 1 refers to the 1988:01-1991:12 subsample; Set 2 refers to the 1992:01-1995:12 subsample. Standard errors are small and omitted.

To sum up, when we combine the results from calls and puts for the vast majority of the options that are actively traded, the delta-hedged gains are overwhelmingly negative. This evidence on the underperformance of delta-hedged portfolios among calls and puts is strongly supportive of a negative volatility risk premium. In a spirit similar to ours, Coval and Shumway (2000) corroborate that (long volatility) at-the-money S&P 500 straddles produce average losses of about 3% per week. That the sign of the market volatility risk premium is negative is in agreement with the parametric approach adopted in Eraker, Johannes, and Polson (2003) and Pan (2002) (i.e., the negative volatility risk premium increases the risk-neutral drift of the volatility process). Exploiting the spanning properties of options, the results in Buraschi and Jackwerth (2001) suggest the possibility of a nonzero volatility risk premium. Although stochastic volatility option models have been shown to reduce fitting errors [Bakshi, Cao, and Chen (1997)], a negative volatility risk premium offers the further potential for reconciling option prices.

Under the premise that higher volatility implies a more negative volatility risk premium, is it empirically true that higher volatility translates into greater underperformance of delta-hedged portfolios? To investigate this issue we constructed the two measures of volatility outlined in Equations (28) and (29), and binned the at-the-money delta-hedged gains into seven volatility groups (<8%, 8-10%, 10-12%, 12-14%, 14-16%, 16-18%, and >18%).

Table 3
Delta-hedged gains for near-the-money calls, by volatility regimes

VOL ^h (%)	T		y ∈ [-2.5%, 0]			y ∈ [0, 2.5%]		
			π (\$)	π/S (%)	π/C (%)	π (\$)	π/S (%)	π/C (%)
<8	428	Mean	0.11	0.04	5.64	-0.25	-0.04	-1.69
		Median	-0.21	-0.04	-7.81	-0.38	-0.08	-3.87
8-10	399	Mean	-0.54	-0.10	-13.52	-0.37	-0.06	-3.03
		Median	-0.72	-0.16	-19.63	-0.57	-0.13	-5.92
10-12	450	Mean	-0.35	-0.07	-6.68	-0.38	-0.08	-3.02
		Median	-0.54	-0.14	-13.18	-0.51	-0.13	-4.92
12-14	295	Mean	-0.44	-0.11	-8.52	-0.22	-0.06	-1.78
		Median	-0.50	-0.14	-11.28	-0.41	-0.11	-4.27
14-16	147	Mean	-1.01	-0.29	-20.90	-1.14	-0.32	-11.97
		Median	-1.12	-0.31	-21.84	-1.13	-0.32	-12.33
16-18	99	Mean	-1.32	-0.41	-24.93	-1.31	-0.40	-12.87
		Median	-1.38	-0.42	-27.36	-1.28	-0.38	-12.34
>18	125	Mean	-1.57	-0.51	-22.90	-1.51	-0.47	-13.41
		Median	-1.42	-0.48	-25.53	-1.62	-0.51	-14.58

Each date is classified into one of seven different volatility regimes based on the annualized estimate of historical volatility (i.e., VOL^h). Delta-hedged gains are computed as described in the note to Table 1. We report (i) the dollar delta-hedged gains, (ii) delta-hedged gains normalized by the index price, and (iii) the delta-hedged gains normalized by the call price, averaged over each volatility regime. Both the mean and the median are displayed. The results are reported for two near-the-money categories: -2.5% to 0% and 0 to 2.5%. The standard errors are small and therefore omitted. T represents the number of days in each volatility classification.

To save on space, we maintain focus on the conventional measure of volatility displayed in Equation (29). The pattern of average delta-hedged gains in Table 3 validates several results of economic relevance. First, consistent with a volatility risk premium, the returns on delta-hedged portfolios are inversely proportional to volatility: during times of higher return volatility, the delta-hedged gains get even more negative. This is true irrespective of whether the underperformance is measured in dollar terms or as a fraction of the index level. Second, the findings are invariant to outliers. In most volatility groups, the median delta-hedged gain is more negative than the mean delta-hedged gain, and generally declines with increases in volatility. Volatility is an important source of the underperformance of delta-hedged portfolios.

4. Delta-Hedged Gains and Option Vega in the Cross Section

We consider next the cross-sectional implication of the volatility risk premium. Following Section 1.2 [Equation (19)], for a fixed σ_t , $\pi_{t,t+\tau}/S_t$ must be related to the option vega, such that mean delta-hedged gains decrease in absolute magnitude for strikes away from at-the-money. We test this implication by adopting the econometric specification

$$\text{GAINS}_t^i = \Psi_0 + \Psi_1 \text{VEGA}_t^i + e_t^i, \quad i = 1, \dots, I, \quad (31)$$

where $\text{GAINS}_t^i \equiv \pi_{t,t+\tau}/S_t$ and VEGA_t^i is the option vega (indexed by monyness $i = 1, \dots, I$). While controlling for volatility and option maturity,

Equation (31) models the proportionality of delta-hedged gains in the option vega. The null hypothesis that volatility risk is not priced corresponds to $\Psi_1 = 0$.

For estimating Equation (31), we require a proxy for VEGA_t^i , and a procedure for controlling for volatility. To demonstrate robustness of the cross-sectional regression estimates, the option vega is approximated in two different ways:

$$\text{VEGA} = \begin{cases} \exp(-d_1^2/2) & \text{Black-Scholes vega,} \\ |y - 1| & \text{Absolute moneyness,} \end{cases} \quad (32)$$

where d_1 is as presented in Equation (30). Two points are worth emphasizing about Equations (31) and (32). First, because $\exp(-d_1^2/2)$ reaches a maximum when the strike is at-the-money, a negative (positive) volatility risk premium corresponds to $\Psi_1 < 0$ ($\Psi_1 > 0$). Furthermore, the magnitude of $\Psi_0 + \Psi_1$ is approximately the mean delta-hedged gains for at-the-money options. Note that the average volatility embedded in d_1 serves simply as a scaling factor for $\log(y)$ and governs the rate of change in $\exp(-d_1^2/2)$, as the option moneyness moves away from the money. For example, for a 30-day option evaluated at 12% volatility, the impact of the risk premium on a 4% away-from-the-money option is half that for at-the-money options. This rate of decrease is slower for higher levels of volatility.

Second, the function $|y - 1|$ reaches a minimum for at-the-money options. In this case the hypothesis of a negative (positive) volatility risk premium corresponds to $\Psi_0 < 0$ and $\Psi_1 > 0$ ($\Psi_0 > 0$ and $\Psi_1 < 0$). In this model the mean delta-hedged gains for at-the-money options is precisely Ψ_0 . Both approximations, $\exp(-d_1^2/2)$ and $|y - 1|$, plausibly characterize the behavior of the option vega.

It is necessary that the sample for each estimation of Equation (31) consists of a panel of delta-hedged gains where the historically measured volatility is approximately constant. To achieve this we divide the sample period into intervals of 2%; within each sample we include all dates where the volatility is within one of these intervals. Therefore we assume the constancy of the volatility risk premium within a volatility classification. To increase the power of the test, and because the sensitivity of the vega (and thus the delta-hedged gains) to moneyness is more pronounced at shorter maturities, we estimate Equation (31) for 30- and 44-day options. With two vega surrogates, we thus have 28 distinct panels, with volatility ranging from approximately 6% to 20%, and with panel size ranging from 46 to 283 observations.

When implementing Equation (31), one econometric issue arises. As there are multiple observations of option prices on each date within a volatility sample, it is possible that there is a date-specific component in $\pi_{i,t+\tau}$ that needs to be explicitly modeled. We follow standard econometric theory [see, e.g., Greene (1997)] and allow for either a date-specific *fixed* effect or a

date-specific *random* effect. In the fixed effects model, we replace Ψ_0 in Equation (31) with $\Psi_{0,t}$. In the random effects model, we allow for a component of the disturbance to be date specific, as modeled by $e_t^i = u_t + v_t^i$. We conduct specification tests on our samples, and in the majority of the samples, the Hausman test of fixed versus random effects and a Lagrange multiplier test of random effects versus ordinary least squares (OLS) favors the random effects specification. As a consequence, all reported results are based on the random effects model, where the coefficients are estimated by feasible generalized least squares panel regression (hereafter FGLS).

Table 4 supports the central implication that the volatility risk premium is negative. Consider first 30-day options and vega measured by $\exp(-d_1^2/2)$. In this case, as hypothesized, the coefficient Ψ_1 is persistently negative. The regression coefficient Ψ_1 ranges between -0.67 and -0.06 , and implies a negative volatility risk premium. For five of seven volatility levels, the coefficient is statistically significant with a minimum (absolute) z -statistic of 2.95 (shown in square brackets). The estimate of $\Psi_0 + \Psi_1$ is roughly in line with the findings in Tables 1 and 3: the mean delta-hedged gains are more negative for higher-volatility versus lower-volatility regimes. For instance, the

Table 4
Delta-hedged gains and option vega: cross-sectional tests

VOL ^h	N	30-day option				N	44-days option			
		Vega is $\exp(-d_1^2/2)$		Vega is $ y - 1 $			Vega is $\exp(-d_1^2/2)$		Vega is $ y - 1 $	
		Ψ_0	Ψ_1	Ψ_0	Ψ_1		Ψ_0	Ψ_1	Ψ_0	Ψ_1
<8	241	0.057 (1.65)	-0.10 [-2.95]	-0.068 [-1.81]	2.68 [7.35]	300	0.067 [1.72]	-0.03 [-0.86]	-0.001 [-0.23]	1.69 [3.49]
8-10	212	0.046 (1.22)	-0.18 [-4.15]	-0.142 [-3.32]	2.73 [4.60]	247	0.195 [2.38]	-0.36 [-10.57]	-0.182 [-2.21]	4.84 [11.17]
10-12	283	0.018 (0.26)	-0.06 [-1.02]	-0.054 [-0.75]	1.02 [1.27]	219	0.070 [0.75]	-0.14 [-2.05]	-0.079 [-0.87]	1.65 [1.98]
12-14	177	0.029 (0.53)	-0.19 [-3.94]	-0.179 [-3.25]	2.44 [4.16]	125	0.062 [0.81]	-0.41 [-8.35]	-0.387 [-5.27]	4.46 [8.52]
14-16	83	0.129 (1.61)	-0.53 [-10.83]	-0.450 [-5.67]	6.00 [10.83]	55	0.025 [0.13]	-0.29 [-1.38]	-0.292 [-2.04]	2.70 [1.34]
16-18	49	0.029 (0.47)	-0.67 [-8.69]	-0.704 [-12.66]	6.78 [8.21]	65	0.283 [1.95]	-0.66 [-5.90]	-0.443 [-3.45]	5.64 [6.01]
>18	46	-0.052 (-0.17)	-0.47 [-1.69]	-0.586 [-2.27]	3.88 [1.79]	78	0.421 [1.96]	-0.86 [-4.15]	-0.522 [-3.59]	5.22 [4.29]

We estimate an FGLS random effect panel regression of delta-hedged gains on the VEGA:

$$\text{GAINS}_t^i = \Psi_0 + \Psi_1 \text{VEGA}_t^i + \epsilon_t^i,$$

$$\epsilon_t^i = u_t + \omega_t^i,$$

where $\text{GAINS}_t^i \equiv \pi_{t,t+\tau}^i / S_t$ (for strike K_i , i, \dots, I). We use two proxies for the option vega. First, VEGA is defined as $\exp(-d_1^2)$, where $d_1 \equiv \frac{1}{\sigma\sqrt{\tau}} \log(y) + \frac{1}{2} \sigma\sqrt{\tau}$, where y_t is option moneyness corresponding to strike K_i . Second, VEGA is the absolute level of moneyness, $|y - 1|$. The data consists of monthly observations of calls over the period January 1988–December 1996. We focus on options with maturity of 30 and 44 days. The sample is chosen such that the volatility is within a predefined interval, VOL^h. N is the number of observations. Since the estimation method is not least squares, the coefficient of determination is omitted. Numbers in square brackets show the z -statistics [Greene (1997)].

estimate of $\Psi_0 + \Psi_1$ is -0.13% in the 8–10% volatility grouping in comparison to -0.41% in the 14–16% volatility grouping. Based on the Wald test, the hypothesis $\Psi_0 + \Psi_1 = 0$ is rejected at the usual significance level (for most groups). Since the R^2 is not particularly instructive for panel regressions, it has been excluded. The results for the 44-day options are comparable with five of seven significantly negative Ψ_1 coefficients. Therefore, for both maturities, the absolute value of delta-hedged gains is maximized for at-the-money options and decrease with the option vega.

When vega is proxied by $|y - 1|$, there is evidence for the joint hypothesis that $\Psi_0 < 0$ and $\Psi_1 > 0$. For 30-day options, Ψ_1 varies from a low of 1.02 to a high of 6.78, and is statistically significant in five of seven estimations. The estimated Ψ_0 coefficient and the associated t -statistics allow us to reject the hypothesis that $E_t(\pi_{t,t+\tau}/S_t) = 0$ (in five of seven volatility groups). As before, the results for 44-day options are consistent with those for 30-day options. Both sets of estimations verify that mean delta-hedged gains decrease in absolute magnitude for strikes away from at-the-money. Our evidence supports the cross-sectional implication of a negative market volatility risk premium.

In Table 5 we provide additional confirmatory evidence for 30-day options. First, in panel A, we report the results from a panel regression when Equation (31) is altered to $\text{GAINS}_t^i = \Psi_0 + \Psi_1 \text{VOL}_t^h \times \text{VEGA}_t^i + e_t^i (i = 1, \dots, I)$. In this specification test we also allow the mean delta-hedged gains to vary with volatility. For example, the time t at-the-money delta-hedged gains are now represented by $\Psi_0 + \Psi_1 \text{VOL}_t$. As observed, the results reported in panel A of Table 5 and those reported in Table 4 are mutually consistent. Second, panel B of Table 5 substantiates that similar results can be found in the subsamples. Therefore our key findings are robust across subsamples and to modifications in the test specifications.

To summarize, the cross-sectional regressions support three main empirical results. The first conclusion that emerges is that we can formally reject the hypothesis that $E_t[\pi_{t,t+\tau}/S_t] = 0$. Moreover, the signs of the estimated coefficients are compatible with a negative volatility risk premium. Finally, the delta-hedged gains are maximized for at-the-money options, and decrease in absolute value for moneyness levels away from at-the-money. Each finding is consistent with the theoretical predictions.

A negative market volatility risk premium has the interpretation that investors are willing to pay a premium to hold options in their portfolio, or that a long position in an index option acts as a hedge to a long position in the market portfolio. We illustrate this point from two different angles. First, we directly examine how option prices react to volatility. For a fixed option maturity, we build a monthly time series of at-the-money call option prices (divided by the index level) and regress it on historical volatility [as

Table 5
Robustness results, delta-hedged gains, and option vega (30-day calls)

Panel A: The specification is $GAINS_t^i = \Psi_0 + \Psi_1 VOL_t^h \times VEGA_t^i + \epsilon_t^i$

VOL ^h (%)	N	VEGA is exp(-d ₁ ²)		VEGA is y-1	
		Ψ ₀	Ψ ₁	Ψ ₀	Ψ ₁
<8	158	0.087 [2.08]	-0.22 [-3.93]	-0.073 [-1.58]	4.30 [5.86]
8-10	212	0.040 [1.04]	-0.18 [-3.77]	-0.136 [-3.30]	2.81 [4.35]
10-12	283	0.26 [0.37]	-0.07 [-1.27]	-0.060 [-0.85]	1.09 [1.48]
12-14	177	0.026 [0.47]	-0.14 [-3.82]	-0.177 [-3.22]	1.86 [4.07]
14-16	83	0.129 [1.60]	-0.36 [-10.97]	-0.450 [-5.70]	4.02 [10.98]
16-18	49	0.024 [0.34]	-0.39 [-8.60]	-0.701 [-13.24]	3.89 [8.18]
>18	44	-0.061 [-0.20]	-0.21 [-1.63]	-0.589 [-2.33]	1.82 [1.85]

Panel B: Subsample results for 1992:01-1995:12

VOL ^h (%)	N	VEGA is exp(-d ₁ ²)		VEGA is y-1	
		Ψ ₀	Ψ ₁	Ψ ₀	Ψ ₁
<8	158	0.084 [1.97]	-0.13 [-3.55]	-0.071 [-1.55]	2.82 [5.60]
8-10	184	0.049 [1.43]	-0.26 [-4.15]	-0.210 [-5.78]	3.73 [8.85]
10-12	148	-0.013 [-0.28]	-0.27 [-7.02]	-0.293 [-6.24]	3.60 [7.01]
12-14	177	0.029 [0.53]	-0.19 [-3.94]	-0.179 [-3.25]	2.44 [4.16]

In panel A, we report the results from the following FGLS random effect panel regression: $GAINS_t^i = \Psi_0 + \Psi_1 VOL_t^h \times VEGA_t^i + \epsilon_t^i$, and $\epsilon_t^i = u_i + \omega_t^i$. The dependent variable is $GAINS_t^i \equiv \pi_{t,t+\tau}^i / S_t^i (i, \dots, I)$. VEGA is defined as either $\exp(-d_1^2)$ or $|y-1|$, where $d_1 = \frac{1}{\sigma\sqrt{\tau}} \log(y) + \frac{1}{2} \sigma\sqrt{\tau}$. As before, y is the option moneyness. The data consist of monthly observations of calls over the period January 1988-December 1995. N is the number of observations. Numbers in square brackets show z -statistics [Greene (1997)]. In panel B we repeat the analysis of Table 4 for the 1992:01-1995:12 subsample. All results are for 30-day calls.

estimated in Equation (29) for $\tau = 30$ days]:

$$\begin{aligned}
 &30 \text{ days: } C_t/S_t \\
 &= 0.004 + 0.05 VOL_t^h + 0.44C_{t-1}/S_{t-1} + e_t, \quad R^2 = 43.16\%, \quad DW = 2.01, \\
 &\quad [3.36] \quad [3.23] \quad [4.12]
 \end{aligned}$$

$$\begin{aligned}
 &44 \text{ days: } C_t/S_t \\
 &= 0.003 + 0.06 VOL_t^h + 0.52C_{t-1}/S_{t-1} + e_t, \quad R^2 = 65.23\%, \quad DW = 2.28. \\
 &\quad [2.26] \quad [3.89] \quad [5.18]
 \end{aligned}$$

Controlling for movements in the index level through time, these regressions show, as would be expected, that call prices respond positively to volatility. To put the estimated slope coefficient in perspective, we note that, in our

sample, the average *C/S* is 1.70% and the average volatility is 11%. An increase in the level of volatility from 11% to 12% will increase *C/S* from 1.70% to 1.79%. This increase is the order of magnitude as that implied by the 30-day at-the-money Black–Scholes vega. Given the extensive evidence on the negative correlation between stock returns and volatility, the positive estimate of the empirical vega confirms the hedging role of options.

To highlight the value of the option as a hedge during significant market declines, we contrast the change in the relative value of index options for the largest 20 negative and positive daily returns (roughly a 3 standard deviation event). On the day prior to a tail event, we buy a nearest-to-the-money short-term call option and compute the Black–Scholes implied volatility. Proceeding to the day after the tail event, we recompute the Black–Scholes implied volatility for the prevailing nearest-to-the-money calls. For each of the largest extreme movements, Table 6 reports (i) the (annualized) implied volatilities and (ii) the corresponding change in implied volatility as a fraction of the implied volatility of the option bought. We can observe that the average change (relative change) in the implied volatility is 1.71% (10.58%) to a downward movement versus -0.84% (-1.53%) to an upward movement.

Table 6
Changes in valuation of index options during large moves

Largest negative price movements					Largest positive price movements				
Date	Price move	Prior IMPL	Subs. IMPL	Relative change	Date	Price move	Prior IMPL	Subs. IMPL	Relative change
880108	-7.00	26.69	28.90	8.26	910117	3.66	27.69	21.29	-23.11
891013	-6.32	13.23	11.87	-10.32	880531	3.41	19.21	19.25	0.22
880414	-4.44	20.05	21.02	4.84	900827	3.18	28.96	24.21	-16.40
911115	-3.73	12.31	15.90	29.16	901001	2.90	23.27	21.61	-7.15
900823	-3.04	25.90	28.96	11.83	910821	2.90	16.04	13.05	-18.67
900806	-3.00	20.84	23.01	10.37	891016	2.75	11.87	20.48	72.60
880120	-2.72	24.14	25.71	6.51	880406	2.66	20.16	19.36	-3.98
901009	-2.70	24.83	24.15	-2.75	910211	2.58	19.07	19.58	2.69
900122	-2.61	17.16	19.61	14.26	911223	2.50	13.99	12.80	-8.49
900112	-2.48	17.00	18.02	6.00	880115	2.49	29.03	25.38	-12.59
930216	-2.41	11.20	15.41	37.57	880728	2.39	17.27	18.79	8.80
910819	-2.37	13.55	16.04	18.42	880608	2.37	20.45	21.17	3.57
890317	-2.27	14.22	15.69	10.30	880902	2.37	19.06	17.40	-8.69
900816	-2.26	19.76	24.11	21.98	900511	2.37	16.60	16.93	1.96
940204	-2.24	9.14	11.10	21.50	880125	2.31	25.10	25.38	1.10
900924	-2.14	24.37	23.70	-2.77	901018	2.31	25.18	24.26	-3.65
881111	-2.13	18.15	18.61	2.50	880729	2.24	17.58	17.48	-0.59
880324	-2.09	21.35	23.69	10.94	890512	2.22	14.30	14.27	-0.16
900821	-2.03	24.21	25.90	6.99	910402	2.19	16.24	16.43	1.22
910510	-1.96	14.02	14.86	6.01	901019	2.18	25.25	20.40	-19.19
Avg.	-3.00	18.61	20.31	10.58	Avg.	2.60	20.32	19.48	-1.53

We proxy valuation changes in index calls by the corresponding change in Black–Scholes implied volatility. This is done in two steps. First, on the day prior to a large daily move, we buy a call option and compute the Black–Scholes implied volatility. Second, proceeding to the day after the large move, we recompute the Black–Scholes implied volatility. We report four sets of numbers: (i) the price movement (in %), (ii) the prior-day implied volatility (denoted as “Prior IMPL”), (iii) the subsequent-day implied volatility (denoted as “Subs. IMPL”), and (iv) the corresponding change in implied volatility as a fraction of the implied volatility of the option bought (i.e., the relative change). The sample period is 1988–1995. In each implied volatility calculation, the index level is adjusted by the present discounted value of dividends. Only short-term call options with strikes that are closest to at-the-money are considered. We display the results from the largest 20 percentage price movements.

Holding everything else constant, the index options become more expensive during stock market declines (in 18 of 20 moves the implied volatility increases). On the other hand, when the market has a strong positive return, the effect on option values is not as striking. These findings further support our assertion that equity index options are desirable hedging instruments.

**5. Delta-Hedged Gains and the Volatility Risk Premium:
Time-Series Evidence**

Following Proposition 2, we now consider the time-series implications of the volatility risk premium for at-the-money options. Fixing option maturity, we estimate the time-series regression

$$\text{GAINS}_t = \Omega_0 + \Omega_1 \text{VOL}_t + \Omega_2 \text{GAINS}_{t-1} + \epsilon_t, \tag{33}$$

where GAINS_t represents the dollar delta-hedged gains for at-the-money options divided by the index level, and VOL_t is the estimate of historical volatility computed over the 30 calendar day period prior to t [see Equation (28) for VOL_t^g , and Equation (29) for VOL_t^h]. In the time-series setting of Equation (33), testing whether volatility risk is not priced is equivalent to testing the null hypothesis $\Omega_1 = 0$. Observe that we have added a lagged value of GAINS to correct for the serial correlation of the residuals. The estimation is done using OLS and the reported t -statistics are based on the Newey–West procedure (with a lag length of 12). As a check, we also estimate the model using the Cochrane–Orcutt procedure for first-order autocorrelation. Since the results are virtually the same, they are omitted to avoid duplication.

To ensure that the regression results are not an artifact of option maturity, we perform regressions at the monthly frequency using delta-hedged gains realized over (i) 30 days, (ii) 44 days, and (iii) 58 days. Although the 30-day series for delta-hedged gains is nonoverlapping, a partial overlap exists with the 44-day and 58-day series. To begin, consider the 30-day series for VOL_t^h . The results in Table 7 show that the OLS estimates of the volatility coefficient, Ω_1 , are negative and statistically significant in all the samples. Over the full sample, the estimated Ω_1 is -0.032 with a t -statistic of -4.39 . In addition, the serial correlation coefficient, Ω_2 , is negative with a t -statistic of -3.47 . The inclusion of GAINS_{t-1} leads to residuals that show little autocorrelation, as is evident from the Box–Pierce statistic with six lags (denoted as Q_6). The coefficient Ω_1 is comparable across maturities and is significantly negative throughout.

The empirical fit of the regressions is reasonable, with the adjusted R^2 higher for each of the two subsamples. Furthermore, the magnitude of the coefficient Ω_0 is an order smaller than that of Ω_1 . Overall, our results seem to indicate that variations in at-the-money delta-hedged gains are related to variations in historical volatility. This result also holds when volatility is

Table 7
Delta-hedged gains and volatility risk premium: time-series regressions

τ	Sample	Historical volatility, VOL_t^h					GARCH volatility, VOL_t^g				
		Ω_0 ($\times 10^{-2}$)	Ω_1	Ω_2	R^2 (%)	Q_6	Ω_0 ($\times 10^{-2}$)	Ω_1	Ω_2	R^2 (%)	Q_6
30 days	Full	0.22	-0.032	-0.199	10.80	1.78	0.05	-0.017	-0.282	6.74	2.34
		[2.00]	[-4.39]	[-3.47]		(0.94)	[0.41]	[-1.71]	[-6.04]		(0.88)
	Set 1	0.87	-0.073	-0.137	15.35	1.10	1.36	-0.101	-0.361	12.55	1.91
		[2.09]	[-2.66]	[-1.36]		(0.98)	[2.04]	[-1.07]	[-7.77]		(0.93)
	Set 2	0.32	-0.051	0.058	14.70	4.45	0.69	-0.089	-0.067	3.81	4.20
		[2.45]	[-3.92]	[0.55]		(0.62)	[1.62]	[-1.97]	[-0.71]		(0.65)
44 days	Full	0.38	-0.045	0.125	13.87	1.68	0.15	-0.023	0.073	0.01	3.37
		[2.53]	[-4.27]	[1.70]		(0.95)	[1.13]	[-2.23]	[1.09]		(0.76)
	Set 1	1.01	-0.080	0.077	24.75	3.06	1.91	-0.135	-0.093	13.00	5.43
		[4.37]	[-5.34]	[0.96]		(0.80)	[4.03]	[-4.49]	[-1.23]		(0.49)
	Set 2	0.52	-0.070	0.422	36.22	6.52	0.99	-0.118	0.317	24.21	9.34
		[4.31]	[-5.23]	[7.04]		(0.37)	[4.40]	[-4.97]	[7.21]		(0.16)
58 days	Full	0.40	-0.048	0.217	11.81	2.81	0.03	-0.013	0.199	2.49	3.09
		[1.92]	[-3.50]	[1.45]		(0.83)	[0.17]	[-0.89]	[1.49]		(0.80)
	Set 1	1.24	-0.099	0.132	19.00	2.31	1.82	-0.131	0.022	14.65	2.98
		[3.07]	[-4.21]	[0.71]		(0.89)	[2.40]	[-2.76]	[0.15]		(0.81)
	Set 2	0.49	-0.066	0.510	36.56	2.18	1.06	-0.126	0.311	24.79	2.30
		[4.69]	[-6.57]	[10.12]		(0.90)	[2.83]	[-3.18]	[9.59]		(0.88)

The regression results are based on the following specification for delta-hedged gains and realized volatility:

$$\text{GAINS}_t = \Omega_0 + \Omega_1 \text{VOL}_t + \Omega_2 \text{GAINS}_{t-1} + e_t$$

where $\text{GAINS}_t \equiv \frac{\pi_{t,t+\tau}}{S_t}$. VOL_t represents the prior month realized volatility. The null hypothesis is that $\Omega_1 = 0$. We include a lagged dependent variable to correct for the serial correlation of the residuals (the Cochrane–Orcutt procedure yields similar inferences). The table reports the coefficient estimate, the t -statistic (in square brackets), the adjusted R^2 , and the Box–Pierce statistic with six lags (denoted Q_6). The p -values for Q_6 are in parentheses. The t -statistics are based on the Newey–West procedure with a lag length of 12. Full refers to the entire sample period of 1988:01–1995:12; Set 1 corresponds to the subsample of 88:1–91:12; and Set 2 corresponds to the subsample of 92:01–95:12. The results are reported for closest to at-the-money calls (with average moneyness of 1.004). For comparison, the regressions are performed using both the historical volatility, VOL_t^h , and the GARCH volatility, VOL_t^g . The GARCH parameters are updated annually, using 1 year of daily return observations. All regressions use call options sampled monthly, with a constant maturity of 30 days, 44 days, and 58 days, respectively.

measured by VOL_t^g . However, the adjusted R^2 's are consistently higher with VOL_t^h . This indicates that a measure of volatility that puts more weight on the recent return history has greater explanatory power and is more informative about delta-hedged portfolio returns. Our repeated finding that $\Omega_1 < 0$ has the implication that the market volatility risk premium is negative.

Is the magnitude of the risk premium indicated by Ω_1 economically significant? Consider again the 30-day series for VOL_t^h . Evaluating Equation (33) at the estimated parameter values, we estimate the effect of the volatility risk premium as measured by the implied dollar delta-hedged gains (at three representative volatility levels):

1. On August 19, 1992, the volatility level was 8.05% with at-the-money call price and index level of \$5.44 and 418.67, respectively. The volatility risk premium is -3.63% of the call option value.
2. Now consider July 19, 1989, where the volatility level was 12.04% with at-the-money call price and index level of \$6.19 and 334.92, respectively. The volatility risk premium is -11.18% of the call option value.

3. Finally, on November 20, 1991, the volatility level was 15.86% with at-the-money call price and index level of \$6.94 and 378.80, respectively. In this case, the volatility risk premium is -19.60% of the value of the call.

Overall, the magnitudes of the volatility risk premium embedded in at-the-money delta-hedged gains are plausible and economically large. The impact of the volatility risk premium is more prominent during times of greater stock market uncertainty. As emphasized in the previous section, this effect may be related to demand for options as hedging instruments.

5.1 Robustness of findings

Several diagnostic tests are performed to examine the stability of Ω_1 . First, we reestimated the regression using the variance as an explanatory variable, with no material change in the results. This last conclusion is not surprising, as the standard deviation and variance are highly correlated. In fact, a model with both variables included performs worse than a model with either of these variables. This suggests that not much can be gained by modeling $\pi_{t,t+\tau}/S_t$ as a polynomial in volatility.

Second, to evaluate whether the results are sensitive to a trending stock market, we reestimated the model using dollar delta-hedged gains, $\pi_{t,t+\tau}$. Again, the results were invariant to this change in specification. Third, we explored the possibility that volatility may be nonstationary. To investigate the impact of nonstationarity on the parameter estimates, we performed an OLS estimation in first differences rather than in levels. This extended specification again points to a negative Ω_1 (these results are available upon request). The principal finding that the market volatility risk premium is negative is robust under alternative specifications.

A natural question that arises is: How sensitive are the results to the mis-measurement of the hedge ratio? Extant theoretical work suggests that the Black–Scholes hedge ratio can depart from the stochastic volatility counterpart when volatility and stock returns are correlated. Guided by this presumption, we now examine (i) whether a negative correlation biases the estimate of $\pi_{t,t+\tau}$, and, if so, (ii) whether our conclusions about the negative volatility risk premium are robust. For each maturity we assemble a time-series of at-the-money calls where the return, $R_{t,t+\tau}$, is positive. For this sample, it is likely that underhedging (overhedging) results in higher (lower) delta-hedged gains.³

³ The logic behind this exercise can be explained as follows. Suppose that the difference between the true hedge ratio and the Black–Scholes hedge ratio is $\xi_t(\sigma_t; y)$, where $\xi > 0$ if Black–Scholes underhedges and negative otherwise. From the definition of delta-hedged gains, it immediately follows that the bias in its estimate is equal to $\Pi_{t,t+\tau} = \int_t^{t+\tau} \xi_u dS_u - \int_t^{t+\tau} r \xi_u S_u du$, which has an expected value of $\int_t^{t+\tau} \bar{\xi}(\mu - r) S_u du$, where μ is the drift of the price process and $\bar{\xi}$ represents the expectation of ξ (assuming ξ is independent of the entire path of S_u). Thus the expected delta-hedged gain is on the order of the market risk premium. If $\mu - r > 0$, and Black–Scholes underhedges the call, then the estimated delta-hedged gain is biased upward.

We estimate the regression $\text{GAINS}_t = \Omega_0 + \Omega_1 \text{VOL}_t^h + \Omega_2 \text{GAINS}_{t-1} + \Omega_3 R_{t,t+\tau} + \epsilon_t$, with the additional variable added to capture the effect of systematic mishedging. In a trending market, we expect $\Omega_3 > 0$, if the call is consistently underhedged, and $\Omega_3 < 0$, if it is overhedged. Although not reported in a table, two findings are worth documenting. First, even when we explicitly account for the impact of under- or overhedging, the coefficient Ω_1 is significantly negative. Second, the coefficient Ω_3 is positive, and hence $\pi_{t,t+\tau}$ is upwardly biased. However, in none of the regressions is Ω_3 statistically significant. That Ω_1 is significantly negative appears robust to errors in hedging arising from a correlation between the stock return process and the volatility process. One interpretation is that the hedge ratio takes into consideration time-varying GARCH volatility and is therefore less misspecified. That a misspecified hedge ratio cannot account for the large negative delta-hedged gains that are observed for at-the-money options is also the conclusion of our simulation results below.

One final cause of concern is that the theoretical distribution of delta-hedged gains may vary across the sample set in a complex manner. Therefore standard procedures adopted in estimating Equation (33) may not fully account for changes in the covariance matrix of $\pi_{t,t+\tau}/S_t$. To explore this, we repeated our estimation using generalized method of moments [Hansen (1982)]. The instrumental variables are a constant and three lags of volatility. For 30-day maturity options, the estimated Ω_1 is -0.045 with a t -statistic of -5.08 (using Newey–West with 12 lags). The minimized value of the GMM criterion function, which is distributed $\chi^2(2)$, has a value of 2.77 and a p -value of 0.24. The results are similar for options of 44 and 58 days. Thus we do not reject the empirical specification in Equation (33). The volatility risk premium coefficient, Ω_1 , is significantly negative, in line with our earlier findings.

5.2 Simulation evidence

Since the empirical tests reject the null hypothesis that volatility risk is unpriced, we pose two additional questions using simulated data: (i) How severe is the small sample bias? and (ii) What is the impact of using Black–Scholes hedge ratio as the approximation for the true hedge ratio? For this artificial economy exercise, our null hypothesis is that volatility is stochastic, but not priced. Therefore we set $\lambda[\sigma_t] = 0$, so that the dynamics of σ_t require no measure change conversions. We simulate the paths of $\{(S_t, \sigma_t) : t = 1, \dots, T\}$, according to Equations (50) and (51). To be consistent with our empirical work, the simulated sample path is taken to be 8 years (2880 days).

At the beginning of the month, an at-the-money call option is bought and delta-hedged discretely over its lifetime. Proceeding to the next month, we repeat this delta-hedging procedure. The option price is given by the stochastic volatility model of Heston (1993). For comparison, the delta-hedged gains

are computed using the hedge ratio from the true stochastic volatility option model as well as using the Black–Scholes model. Across each simulation run we generate 96 observations on delta-hedged gains and the prior 30-days volatility. Using the simulated sample, we estimate Equation (33). In Table 8 we report the sample distribution of estimated coefficients over 1000 simulations for two option maturities, 30 days and 44 days (the mean, and the mean absolute deviation in curly brackets). The first point to note is that with unpriced volatility risk, the mean delta-hedged gains are virtually zero. Under the stochastic volatility model, the magnitude of π/S is several orders lower than those depicted in Table 1. Second, the use of Black–Scholes delta imparts a negligible bias. For example, for 30-day options, the mean π/S is -0.0018% with the stochastic volatility hedge ratio versus 0.0022% with the Black–Scholes hedge ratio. In conclusion, the simulations show that the use of the Black–Scholes hedge ratio does not perversely bias the magnitude of delta-hedged gains.

Now shift attention to the sample rejection level of the estimated coefficients from the simulated data. First, given the theoretical p -value benchmark of 5%, the null hypothesis $\Omega_1 = 0$ should be rejected only occasionally. Again consider the stochastic volatility model with option maturity of 30 days. Inspection of Table 8 shows that the frequency of $t(\Omega_1) < -2$ is 3.94%. Moreover, the frequency of $t(\Omega_1) > 2$ is 1.21%. Therefore, when combined, there is only a small overrejection of the null hypothesis. If the hedge ratio is

Table 8
Properties of delta-hedged gains in simulated economies

τ	Hedge ratio	Simulated delta-hedged gains			Simulated coefficient values $GAINS_t = \Omega_0 + \Omega_1 VOL_t + \Omega_2 GAINS_{t-1} + \epsilon_t$					
		π (\\$)	π/S (%)	π/C (%)	Ω_0	Ω_1	Ω_2	R^2	Freq. for $t(\Omega_1) > 2$	Freq. for $t(\Omega_1) < -2$
30 days	SV	-0.0024 {0.0235}	-0.0018 {0.0015}	-0.1641 {0.1358}	0.046 {0.170}	-0.519 {1.412}	-0.019 {0.093}	2.43% {1.74}	0.0121	0.0394
	BS	0.0036 {0.0233}	0.0022 {0.0015}	0.2039 {0.1346}	0.055 {0.171}	-0.555 {1.407}	-0.020 {0.093}	2.44% {1.73}	0.0120	0.0420
44 days	SV	0.0022 {0.0305}	0.0017 {0.0021}	0.1372 {0.1629}	0.077 {0.137}	-0.616 {0.922}	0.149 {0.092}	4.90% {3.17}	0.0050	0.0710
	BS	0.0089 {0.0297}	0.0021 {0.0015}	0.5151 {0.1576}	0.077 {0.136}	-0.615 {0.923}	0.149 {0.093}	4.90% {3.17}	0.0050	0.0710

We simulate delta-hedged gains in an economy where volatility is stochastic, but not priced. The simulation experiment (see Appendix B for more details) is based on the following discretization of stock returns and volatility: $S_{t+h} = S_t + \mu S_t h + \sigma_t S_t \epsilon_t^1 \sqrt{h}$ and $\sigma_{t+h}^2 = \sigma_t^2 + \kappa(\theta - \sigma_t^2)h + v \sigma_t \epsilon_t^2 \sqrt{h}$, where h is set equal to 1 day. For the simulation $S_0 = 100$, $\sigma_0 = 10\%$, $\kappa = 2$, $\theta = 0.01$, $\rho = -0.5$, and $v = 0.1$. In each simulation run, the theoretical call option values are generated according to the stochastic volatility option pricing model of Heston (1993). Throughout we assume that the interest rate and dividend yield are zero, and $\lambda[\sigma_t] = 0$. The delta-hedged gains are computed in two different ways. First, we use the hedge ratio given by the stochastic volatility model (denoted as “SV”), and second using the Black–Scholes model (denoted as “BS”). We consider at-the-money options with a maturity of 30 days (and 44 days). Every 30 (44) days, the option is bought and delta hedged. Over 8 years (2880 days), this produces 96 monthly observations on delta-hedged gains and prior 30 days volatility. For each simulation, we perform the time-series regression $GAINS_t = \Omega_0 + \Omega_1 VOL_t + \Omega_2 GAINS_{t-1} + \epsilon_t$. The reported Ω_0 , Ω_1 , Ω_2 , and adjusted R^2 are averages over 1000 simulations. The mean absolute deviation of the estimate is shown in curly brackets. We report the frequency of significant Ω_1 (i.e., $t(\Omega_1) > 2$ and $t(\Omega_1) < -2$). We also show the mean delta-hedged gains across all simulations for both the SV model and the BS model.

replaced with the BS delta, the simulated rejection frequency is again close to the theoretical 5%.

Because the 44-day options allow for some overlap in the data, we expect to see autocorrelation and therefore poorer small sample properties. The simulations confirm that the frequency of the rejection of the hypothesis of $\Omega_1 = 0$ is slightly higher, at 7.1%. However, the overlap does not affect the estimate of the mean π/S (which is 0.0017%); neither does it worsen the fit with the Black–Scholes hedge ratio. Overall the simulation evidence suggests that small sample biases are not large, and that the use of the Black–Scholes hedge ratio has a negligible effect on the estimations. Having said this, we can now proceed to examine the jump-fear foundations of negative delta-hedged portfolio returns.

6. Delta-Hedged Gains and Jump Exposures

While the body of evidence presented so far appears consistent with a volatility risk premium, the losses on the delta-hedged portfolios may also be reconciled by the fear of stock market crashes. The underlying motivation is that option prices not only reflect the physical volatility process and the volatility risk premium, but also the potential for unforeseen tail events [Jackwerth and Rubinstein (1996)]. Jump fears can therefore dichotomize the risk-neutral index distribution from the physical index distribution, even in the absence of a volatility risk premium. Indeed, empirical evidence indicates that the risk-neutral index distribution is (i) more volatile, (ii) more left skewed, and (iii) more leptokurtotic relative to the physical index distribution [Rubinstein (1994), Jackwerth (2000), and Bakshi, Kapadia, and Madan (2003)]. As our characterization of delta-hedged gains shows in Equation (24), these distributional features can induce underperformance of the delta-hedged option strategies. If, in addition, the jump risk premium surfaces more prominently during volatile markets [Bates (2000), Pan (2002), Eraker, Johannes, and Polson (2003)], then it can account for the accompanying greater delta-hedged losses.

To empirically distinguish between the effects of stochastic volatility and jumps on delta-hedged gains, two decisions are made at the outset. One, in the tradition of Jackwerth and Rubinstein (1996), Bates (2000), and Bakshi, Kapadia, and Madan (2003), we assume that jump fears can be surrogated through the skewness and kurtosis of the risk-neutral index distribution. In the modeling framework of Equation (23), the mean jump size governs the risk-neutral skew, and the jump intensity is linked to kurtosis. For instance, the fear of market crashes can impart a left skew and shift more probability mass toward low-probability events. Two, the risk-neutral skewness and kurtosis are recovered using the model-free approach of Bakshi, Kapadia, and Madan (2003). They show that the higher-order risk-neutral moments can be spanned and priced using a positioning in out-of-the-money calls and puts.

In what follows, the relative impact of jump fears on delta-hedged gains is gauged from three perspectives.

First, we modify the time-series specification of Equation (33) to include a role for risk-neutral skew and kurtosis, as shown below:

$$\begin{aligned} \text{GAINS}_t = & \Omega_0^* + \Omega_1^* \text{VOL}_t^h + \Omega_2^* \text{GAINS}_{t-1} + \Omega_3^* \text{SKEW}_t^* \\ & + \Omega_4^* \text{KURT}_t^* + \epsilon_t^*, \end{aligned} \tag{34}$$

where VOL_t^h is the historical volatility, SKEW_t^* is the risk-neutral index skewness, and KURT_t^* is the risk-neutral index kurtosis. For convenience, the exact expressions for skew and kurtosis are displayed in Equations (45) and (46) of the appendix. Specifically the risk-neutral skewness and kurtosis reflect the price of the cubic contract and the kurtic contract, respectively. As before, we include a lagged value of delta-hedged gains to correct for serially correlated residuals. The estimation is by OLS and the t -statistics are computed using the Newey–West procedure with 12 lags. The main idea behind the empirical specification of Equation (34) is to investigate whether physical volatility loses its significance in the presence of such jump-fear proxies as risk-neutral skews and kurtosis. We also employed the slope of the volatility smile and the Bates skewness premium measure as alternative proxies for jump fear and obtained similar conclusions (details are available from the authors). To maintain the scope of the investigation, these extended measures are excluded from the main body of the article.

Before we discuss the estimation results presented in Table 9, it must be stressed that there is substantial evidence of jump fear in the postcrash risk-neutral distributions. Over the entire sample period, the average risk-neutral skewness is -1.38 and the risk-neutral kurtosis is 7.86 for 30-day distributions. These numbers are roughly comparable to those reported in Jackwerth and Rubinstein (1996) for longer-term options and in Bakshi, Kapadia, and Madan (2003) for the S&P 100 index options. The most important point that emerges from Table 9 is that historical volatility continues to significantly affect variations in delta-hedged gains. The coefficient on volatility ranges between -0.111 and -0.041 , and is statistically significant in all nine estimations. The evidence on the role of skew and kurtosis is less conclusive. Although skew enters the regression with the correct sign, it is only marginally significant. The positive estimate of Ω_3^* indicates that a more negatively skewed risk-neutral distribution makes delta-hedged gains more negative from one month to the next. In addition, the sign of kurtosis is contrary to what one might expect. While not reported, skewness (kurtosis) is not individually significant when volatility and kurtosis (skewness) are omitted as explanatory variables in Equation (34). Finally, comparing the empirical fit between Tables 7 and 9, the inclusion of skew and kurtosis only modestly improves the adjusted R^2 (by about 3%). In summary, this exercise suggests

Table 9
Effect of jumps on delta-hedged gains

τ	Sample	Ω_0^* ($\times 10^{-2}$)	Ω_1^*	Ω_2^*	Ω_3^* ($\times 10^{-2}$)	Ω_4^* ($\times 10^{-3}$)	R^2	Q_6
30 days	Full	0.51	-0.041	-0.16	0.31	0.28	13.98	1.26 (0.97)
		[2.86]	[-3.44]	[-2.84]	[1.82]	[1.73]		
	Set 1	1.18	-0.082	-0.10	0.21	0.13	15.46	1.27 (0.97)
		[1.97]	[-1.99]	[-0.94]	[0.90]	[0.58]		
	Set 2	0.63	-0.064	0.04	0.17	0.05	17.31	4.52 (0.61)
		[3.55]	[-4.74]	[0.47]	[1.40]	[0.28]		
44 days	Full	0.45	-0.046	0.19	0.23	0.32	17.26	2.41 (0.91)
		[2.03]	[-4.09]	[3.31]	[1.51]	[2.24]		
	Set 1	1.11	-0.085	0.118	0.04	0.05	23.47	3.36 (0.76)
		[2.84]	[-3.70]	[1.43]	[0.23]	[0.28]		
	Set 2	0.55	-0.068	0.42	0.12	0.17	34.13	6.41 (0.38)
		[3.34]	[-3.95]	[7.01]	[0.70]	[0.59]		
58 days	Full	0.60	-0.055	0.22	0.34	0.41	12.42	1.82 (0.94)
		[1.85]	[-3.08]	[1.30]	[1.46]	[1.79]		
	Set 1	1.71	-0.111	0.10	0.50	0.44	19.03	2.55 (0.86)
		[3.30]	[-3.35]	[0.53]	[3.14]	[4.25]		
	Set 2	0.46	-0.064	0.50	-0.05	-0.08	35.41	3.71 (0.72)
		[3.88]	[-6.88]	[7.04]	[-0.43]	[-0.55]		

We employ skewness and kurtosis of the risk-neutral distribution as proxies for jump fear. The regression results are based on the following specification between delta-hedged gains, historical volatility, and the higher-order moments of the risk-neutral return distribution:

$$\text{GAINS}_t = \Omega_0^* + \Omega_1^* \text{VOL}_t^h + \Omega_2^* \text{GAINS}_{t-1} + \Omega_3^* \text{SKEW}_t^* + \Omega_4^* \text{KURT}_t^* + e_t,$$

where $\text{GAINS}_t \equiv \frac{\pi_{t,t+\tau}}{S_t}$, VOL_t^h represents the historical volatility. To correct for the serial correlation of the residuals, we have included a lagged dependent variable (the Cochrane–Orcutt procedure yields similar inferences). We record the coefficient estimate, the t -statistic (in square brackets), the adjusted R^2 , and the Box–Pierce statistic with six lags (denoted Q_6). The p -values for Q_6 are in parentheses. The t -statistics are based on the Newey–West procedure with a lag length of 12. Full refers to the entire sample period of 1988:01–1995:12; Set 1 corresponds to the subsample of 88:1–91:12; and Set 2 corresponds to the subsample of 92:01–95:12. The results are reported for closest to at-the-money calls. All regressions use call options sampled monthly, with a constant maturity of 30 days, 44 days, and 58 days, respectively. The model-free estimate of risk-neutral skewness, SKEW_t^* , and the risk-neutral kurtosis, KURT_t^* , are constructed as described in the appendix.

that volatility may be of first-order importance in explaining negative delta-hedged gains.

In the second exercise, we study the behavior of delta-hedged option portfolios for a holdout sample when jump fears are much less pronounced. For this purpose we selected the six-month interval from January 1987 through June 1987 (option data provided by Bent Christensen). What is unique about this precrash period is that risk-neutral index distributions are essentially log normal. Especially suited for the task at hand, the jump fears are virtually lacking during this precrash sample [Jackwerth and Rubinstein (1996)]. Table 10 reports the mean delta-hedged gains for out-of-the-money calls and puts. The average delta-hedged gains for near-the-money 14–30 day calls (puts) is \$-0.65 (-\$0.82). In fact, the delta-hedged gains are strongly negative in all the 16 moneyness and maturity categories. Furthermore, the majority of the options have $\pi < 0$, as seen by the large $1_{\pi < 0}$ statistics. The delta-hedged gains are negative in both the precrash and postcrash periods. While not displayed, the average implied volatility for at-the-money options is higher than the historically realized volatility, suggesting that the well-known bias

Table 10
Delta-hedged gains for a precrash period (January 1987–June 1987, calls and puts)

	Moneyiness $y - 1$	π (in \$)			π/S (in %)			$1_{\pi < 0}$ {N}
		14–30	31–60	All	14–30	31–60	All	
Calls	-10% to -7.5%	-0.12 (0.08)	-1.16 (0.23)	-0.73 (0.18)	-0.04 (0.03)	-0.40 (0.08)	-0.26 (0.06)	86% {21}
	-7.5% to -5%	-0.28 (0.10)	-0.77 (0.19)	-0.61 (0.14)	-0.10 (0.04)	-0.27 (0.07)	-0.22 (0.05)	69% {85}
	-5% to -2.5%	-0.64 (0.11)	-1.11 (0.16)	-0.96 (0.11)	-0.22 (0.04)	-0.38 (0.06)	-0.33 (0.04)	73% {238}
	-2.5% to 0%	-0.65 (0.12)	-0.81 (0.14)	-0.75 (0.10)	-0.22 (0.04)	-0.28 (0.05)	-0.25 (0.03)	67% {276}
Puts	0 to 2.5%	-0.82 (0.15)	-0.87 (0.14)	-0.85 (0.10)	-0.29 (0.05)	-0.31 (0.05)	-0.30 (0.03)	74% {282}
	2.5 to 5%	-0.63 (0.11)	-1.03 (0.10)	-0.87 (0.08)	-0.22 (0.04)	-0.36 (0.03)	-0.31 (0.03)	84% {274}
	5 to 7.5%	-0.23 (0.09)	-0.74 (0.07)	-0.53 (0.06)	-0.09 (0.03)	-0.26 (0.02)	-0.19 (0.02)	84% {245}
	7.5 to 10%	-0.07 (0.07)	-0.67 (0.04)	-0.44 (0.04)	-0.03 (0.02)	-0.24 (0.01)	-0.16 (0.01)	92% {150}
All	-0.25	-0.44	-0.37	-0.09	-0.15	-0.13		

For a selected precrash period, we display the delta-hedged gains for out-of-the-money calls (i.e., $y < 1$) and out-of-the-money puts (i.e., $y > 1$). In the case of calls, the discretely rebalanced delta-hedged gains, $\pi_{t, t+\tau}$, is computed as $\pi_{t, t+\tau} = \max(S_{t+\tau} - K, 0) - C_t - \sum_{n=0}^{N-1} \Delta_n (S_{n+1} - S_n) - \sum_{n=0}^{N-1} r_n (C_t - \Delta_n S_n) \frac{\tau}{N}$, where the interest rate, r_n , and the option delta, Δ_n , are updated on a daily basis. The option delta is computed as the Black–Scholes hedge ratio evaluated at the GARCH volatility. The rebalancing frequency, τ/N , is set to 1 day. Reported in the table are (i) dollar delta-hedged gains ($\pi_{t, t+\tau}$) and the (ii) delta-hedged gains normalized by the index level ($\pi_{t, t+\tau}/S_t$). All delta-hedged gains are averaged over their respective moneyiness and maturity category. The standard errors, shown in parentheses, are computed as the sample standard deviation divided by the square root of the number of observations. $1_{\pi < 0}$ is the proportion of delta-hedged gains with $\pi < 0$, and N is the number of options (in curly brackets). “All” combines the delta-hedged gains from maturities of 14–30 days and 31–60 days.

between the implied and the realized volatility predates crash fears and option skews.

In the final evaluation exercise, we compute the average delta-hedged gains for the largest downward and upward market movements. Intuitively, if fears of negative jumps are the predominant driving factor in determining negative delta-hedged gains, then there should be a strong asymmetric effect [see Equation (24)], with large positive index returns not necessarily resulting in large negative delta-hedged gains. In contrast, it may be argued that a negative volatility risk premium would cause large negative delta-hedged gains, irrespective of the sign of the market return. To briefly examine this reasoning, consider closest to at-the-money short-term calls bought on the day subsequent to a tail event. Respectively for the largest 10 (20) tail events, the average scaled delta-hedged gain, $\pi_{t, t+\tau}/S_t$, is -0.52% (-0.43%) on positive return dates, compared with -0.86% (-0.51%) on negative return dates. The evidence indicates that delta-hedged gains become more negative for *both* extreme negative and positive returns. This evidence from the extremes is largely consistent with the regression results.

One overall conclusion that can be drawn is that priced volatility risk is a more plausible characterization for negative delta-hedged gains. It is possible that if some extremely low-probability event is included, the resulting

large positive gain may wipe out all cumulative losses. However, this low-probability event (of the required magnitude) has not yet occurred in our sample. Our key finding that the market volatility risk premium is, on average, negative is mutually consistent with other evidence reported in Benzoni (1999), Poteshman (1998), Jones (2002), and Pan (2002).

7. Final Remarks

Is volatility risk premium negative in equity index option markets? We argue that the central implication of a nonzero volatility risk premium is that the gains on a delta-neutral strategy that buys calls and hedges with the underlying stock are nonzero and determined jointly by the volatility risk premium and the option vega. Specifically we establish that the volatility risk premium and the mean discrete delta-hedged gains share the same sign. It is shown that this implication can be tested by relatively robust econometric specifications in either the cross section of option strikes or in the time series. These tests do not require the identification of the pricing kernel, nor the correct specification of the volatility process.

Using S&P 500 index options, our empirical results indicate that the delta-hedged gains are nonzero, and consistent with a nonzero volatility risk premium. The main findings of our investigation are summarized below:

1. The delta-hedged call portfolios statistically underperform zero (across most moneyness and maturity classifications). The losses are generally most pronounced for at-the-money options.
2. The underperformance is economically significant and robust. When out-of-the-money put options are delta-hedged, a similar pattern is documented. The documented underperformance of delta-hedged option portfolios is consistent with a negative volatility risk premium.
3. Controlling for volatility, the cross-sectional regression specifications provide support for the prediction that the absolute value of delta-hedged gains are maximized for at-the-money options, and decrease for out-of-the-money and in-the-money options.
4. During periods of higher volatility, the underperformance of the delta-hedged portfolios worsens. As suggested by the hypothesis of a negative volatility risk premium, time variation in delta-hedged gains of at-the-money options are negatively correlated with historical volatility. This finding is robust across subsamples, and to mismeasurement of the hedge ratio.
5. Finally, volatility significantly affects delta-hedged gains even after accounting for jump fears. Jump risk cannot fully explain the losses on the delta-hedged option portfolios.

In economic terms, a negative volatility risk premium suggests an equilibrium where equity index options act as a hedge to the market portfolio, and

is consistent with prevailing evidence that equity prices react negatively to positive volatility shocks. Thus investors would be willing to pay a premium to hold options in their portfolio, and this would make the option price higher than its price when volatility is not priced. The empirical results of this article strengthen the view that equity index options hedge downside risk.

There are two natural extensions to this article. First, given that volatilities of individual stocks and the market index comove highly, one could examine whether the volatility risk premium is negative in individual equity options. The cross-sectional restrictions on delta-hedged gains and the volatility risk premium can be tested in the cross section of individual equity options. Second, volatility risk is of importance in almost every market. The analysis conducted here can be directly applied to include other markets such as foreign exchange and commodities. Much more remains to be learned about how volatility risk is priced in financial markets.

Appendix A: Proof of Results

Proof of Proposition 1. We need to show that $E_t(\pi_{t,t+\tau}) = O(1/N)$, when $\lambda_t[\sigma_t] \equiv 0$. First, without loss of generality, assume $r = 0$, $\Delta_t \equiv \partial C_t / \partial S_t$, and $v_t \equiv \eta_t \partial C_t / \partial \sigma_t$. Second, let the period corresponding to the time to expiration, $t = 0$ to $t = t + \tau$, be divided equally into N periods, corresponding to dates, t_n , $n = 0, 1, \dots, N$, where $t_0 = 0$, $t_N = t + \tau$, and $t_n - t_{n-1} = h$.

Consider the delta-hedged gains over one period, from t_n to t_{n+1} . If the volatility risk premium is zero, then from Equation (13),

$$C_{n+1} = C_n + \int_{t_n}^{t_{n+1}} \Delta_u dS_u + \int_{t_n}^{t_{n+1}} v_u dW_u^2, \tag{35}$$

where, for brevity, we intend n to mean t_n . Define the operators, $\mathcal{L}[\cdot] = \frac{\partial}{\partial t}[\cdot] + \mu_t S_t \frac{\partial}{\partial S_t}[\cdot] + \theta_t \frac{\partial}{\partial \sigma_t}[\cdot] + \frac{1}{2} \sigma_t^2 S_t^2 \frac{\partial^2}{\partial S_t^2}[\cdot] + \frac{1}{2} \eta_t^2 \frac{\partial^2}{\partial \sigma_t^2}[\cdot] + \sigma_t S_t \eta_t \frac{\partial^2}{\partial S_t \partial \sigma_t}[\cdot]$, $\Gamma_1 = \sigma_t S_t \frac{\partial}{\partial S_t}[\cdot]$, and $\Gamma_2 = \eta_t \frac{\partial}{\partial \sigma_t}[\cdot]$. Appealing to an Itô–Taylor expansion,

$$\begin{aligned} C_{n+1} = & C_n + \int_{t_n}^{t_{n+1}} \left(\Delta_n + \int_{t_n}^{t_{n+1}} \mathcal{L}[\Delta_t] dt + \int_{t_n}^{t_{n+1}} \Gamma_1[\Delta_t] dS_t + \int_{t_n}^{t_{n+1}} \Gamma_2[\Delta_t] d\sigma_t \right) dS_u \\ & + \int_{t_n}^{t_{n+1}} [v_n + \int_{t_n}^{t_{n+1}} \mathcal{L}[v_t] dt + \int_{t_n}^{t_{n+1}} \Gamma_1[v_t] dS_t + \int_{t_n}^{t_{n+1}} \Gamma_2[v_t] d\sigma_t] dW_u^2. \end{aligned} \tag{36}$$

With an additional Itô–Taylor expansion to include all terms up to $O(h)$, we can rewrite this as,

$$\begin{aligned} C_{n+1} = & C_n + \Delta_n \int_{t_n}^{t_{n+1}} dS_u + \sigma_n^2 S_n^2 \Gamma_1[\Delta_n] \int_{t_n}^{t_{n+1}} \int_{t_n}^{t_{n+1}} dW_t^1 dW_u^1 + v_n \int_{t_n}^{t_{n+1}} dW_u^2 \\ & + \sigma_n S_n \eta_n \Gamma_2[\Delta_n] \int_{t_n}^{t_{n+1}} \int_{t_n}^{t_{n+1}} dW_t^2 dW_u^1 + \eta_n^2 \Gamma_1[v_n] \int_{t_n}^{t_{n+1}} \int_{t_n}^{t_{n+1}} dW_t^1 dW_u^2 \\ & + \sigma_n S_n \eta_n \Gamma_2[v_n] \int_{t_n}^{t_{n+1}} \int_{t_n}^{t_{n+1}} dW_t^1 dW_u^2 + A_0, \end{aligned} \tag{37}$$

where A_0 consists of terms such as $\int_{t_n}^{t_{n+1}} \int_{t_n}^{t_{n+1}} g(S_t, \sigma_t, t) dt du$ and $\int_{t_n}^{t_{n+1}} \int_{t_n}^{t_{n+1}} h(S_t, \sigma_t, t) dW_t^j ds$, $j = S, \sigma$. Under generally accepted regularity conditions [Lemma 2.2 of Milstein (1995)], $E(A_0) = O(h^2)$ and $E(A_0^2) = O(h^3)$. It follows from Theorem 1.1 in Milstein (1995) that the

order of accuracy of the above discretization over the N steps in the interval, $t = 0$ to $t = t + \tau$, is $h = \tau/N$, so that it is of $O(1/N)$. Rearranging Equation (37), we can write $\pi_{t,t+\tau}$ as,

$$\begin{aligned} \pi_{t,t+\tau} &\equiv \sum_{n=0}^{N-1} C_{n+1} - C_n - \Delta_n(S_{n+1} - S_n), \\ &= \sum_{n=0}^{N-1} \left[\sigma_{t_n}^2 S_{t_n}^2 \Gamma_1[\Delta_{t_n}] \int_{t_n}^{t_{n+1}} \int_{t_n}^u dW_t^1 dW_u^1 + v_{t_n} \int_{t_n}^{t_{n+1}} dW_t^2 \right. \\ &\quad + \sigma_{t_n} S_{t_n} \eta_{t_n} \Gamma_2[\Delta_{t_n}] \int_{t_n}^{t_{n+1}} \int_{t_n}^u dW_t^2 dW_u^1 + \eta_{t_n}^2 \Gamma_1[v_{t_n}] \int_{t_n}^{t_{n+1}} \int_{t_n}^u dW_t^2 dW_s^2 \\ &\quad \left. + \sigma_{t_n} S_{t_n} \eta_{t_n} \Gamma_2[v_{t_n}] \int_{t_n}^{t_{n+1}} \int_{t_n}^s dW_t^1 dW_s^2 \right] + O(1/N). \end{aligned} \quad (38)$$

As the expected value of the Itô integrals is zero, the proposition is proved. \blacksquare

Proof of Lemma 1. The proof is by induction. To fix ideas, we prove the case where $\alpha_t(\tau)$. The extension to $\alpha_t(\sigma_t, \tau)$ is straightforward. Consider $\mathcal{L}[\alpha_t(\tau) S^\phi]$,

$$\mathcal{L}[\alpha_t(\tau) S^\phi] = -\frac{\partial \alpha}{\partial \tau} S^\phi + \alpha \mu_t S \frac{\partial S^\phi}{\partial S} + \frac{1}{2} \alpha \sigma_t^2 S^2 \frac{\partial^2 S^\phi}{\partial S^2} \quad (39)$$

$$= \left(-\partial \alpha_t / \partial \tau + \phi \alpha_t(\tau) \mu_t + \frac{1}{2} \phi(\phi - 1) \alpha_t(\tau) \sigma_t^2 \right) S^\phi \quad (40)$$

by assuming $dS_t = \mu_t S_t dt + \sigma_t S_t dW$. This implies that $\mathcal{L}[\mathcal{L}[\alpha_t(\tau) S^\phi]]$ is again proportional to S^ϕ . By induction, $\mathcal{L}^n[g]$, for any $n \in \{1, 2, 3, \dots\}$, is proportional to S^ϕ . \blacksquare

Proof of Proposition 2. The proof relies on evaluating each term in the expansion of Equation (18). We have $E(\Pi_{t,t+\tau}) = g(S_t, \sigma_t) \tau + \frac{1}{2} \tau^2 \mathcal{L}[g(S_t, \sigma_t)] + \frac{1}{6} \tau^3 \mathcal{L}^2[g(S_t, \sigma_t)] + \dots$, where $g(S_t, \sigma_t) = \lambda_t \partial C_t / \partial \sigma_t$. Here the vega is proportional to S_t , and so $\partial C_t / \partial \sigma_t = \beta_t(\tau; y) S_t$.

Under the maintained assumption that $\lambda_t = \lambda \sigma_t$, $g(S_t, \sigma_t) = \lambda \beta_t(\tau; y) S_t \sigma_t$. From that,

$$\mathcal{L}^1[g] = -\lambda(\partial \beta_t / \partial \tau) S_t \sigma_t + \lambda \beta_t \mu_t S_t (\partial S_t / \partial S_t) \sigma_t + \lambda \beta_t S_t (-\kappa \sigma_t) (\partial \sigma_t / \partial \sigma_t), \quad (41)$$

$$= \lambda \varphi_t S_t \sigma_t, \quad (42)$$

where $\varphi_t \equiv -\partial \beta_t / \partial \tau + \beta_t \mu_t - \beta_t \kappa$. From Lemma 1, successive $\mathcal{L}^n[g]$ inherit the same form as Equation (42), as in $\lambda \varphi_n S_t \sigma_t$. Therefore, $E(\Pi_{t,t+\tau}) = \lambda \varphi_t(\tau) S_t \sigma_t$, where $\varphi_t(\tau) \equiv \sum_{n=0}^{\infty} \frac{\tau^{1+n}}{(1+n)!} \varphi_n$. \blacksquare

Proof of Equation (24). Using Equation (23) and applying Itô's lemma, the call option satisfies the dynamics

$$\begin{aligned} C_{t+\tau} &= C_t + \int_t^{t+\tau} \frac{\partial C_u}{\partial S_u} dS_u + \int_t^{t+\tau} \frac{\partial C_u}{\partial \sigma_u} d\sigma_u + \int_t^{t+\tau} b_u du \\ &\quad + \Lambda_J \int_t^{t+\tau} \sigma_u \int_{-\infty}^{\infty} (C(S_u e^x) - C(S_u)) q[x] dx du, \end{aligned} \quad (43)$$

where $C(S_u e^x)$ implies that the option price is evaluated at $S_u e^x$. In Equation (43), $q[x]$ is the physical density of the jump size, x , and $b_u \equiv \frac{\partial C_u}{\partial u} + \frac{1}{2} \sigma_u^2 S_u^2 \frac{\partial^2 C_u}{\partial S_u^2} + \frac{1}{2} \eta_u \frac{\partial^2 C_u}{\partial \sigma_u^2} + \rho \eta_u \sigma_u S_u \frac{\partial^2 C_u}{\partial S_u \partial \sigma_u}$. The call price is a solution to the partial integro-differential equation,

$$\begin{aligned} \frac{1}{2} \sigma_t^2 S^2 \frac{\partial^2 C}{\partial S^2} + \frac{1}{2} \eta_t^2 \frac{\partial^2 C}{\partial \sigma^2} + \rho \eta_t \sigma_t S \frac{\partial^2 C}{\partial S \partial \sigma} + (r - \mu_J^* \Lambda_J \sigma_t) S \frac{\partial C}{\partial S} + (\theta_t - \lambda_t[\sigma]) \frac{\partial C}{\partial \sigma} \\ + \frac{\partial C}{\partial t} - r C + \Lambda_J \sigma_t \int_{-\infty}^{\infty} (C(S_u e^x) - C(S_u)) q^*[x] dx = 0, \end{aligned} \quad (44)$$

for risk-neutral density $q^*[x]$. Combining Equations (43) and (44) and using the definition of $\Pi_{t,t+\tau}$, we get Equation (24). ■

Expressions for risk-neutral skew and kurtosis used in Section 6

The model-free estimates of risk-neutral return skewness and kurtosis are based on Bakshi, Kapadia, and Madan (2002). Specifically, the risk-neutral skewness, $SKEW^*(t, \tau)$, is given by

$$SKEW^*(t, \tau) \equiv \frac{E_t^* \{ (R_{t,t+\tau} - E_t^*[R_{t,t+\tau}])^3 \}}{\{ E_t^* (R_{t,t+\tau} - E_t^*[R_{t,t+\tau}])^2 \}^{3/2}} = \frac{e^{r\tau} W(t, \tau) - 3\mu(t, \tau)e^{r\tau} V(t, \tau) + 2\mu(t, \tau)^3}{[e^{r\tau} V(t, \tau) - \mu(t, \tau)^2]^{3/2}} \tag{45}$$

and the risk-neutral kurtosis, denoted $KURT^*(t, \tau)$, is

$$KURT^*(t, \tau) = \frac{e^{r\tau} X(t, \tau) - 4\mu(t, \tau)e^{r\tau} W(t, \tau) + 6e^{r\tau}\mu(t, \tau)^2 V(t, \tau) - 3\mu(t, \tau)^4}{[e^{r\tau} V(t, \tau) - \mu(t, \tau)^2]^2}, \tag{46}$$

where

$$V(t, \tau) = \int_{S_t}^{\infty} \frac{2(1 - \ln[\frac{K}{S_t}])}{K^2} C(t, \tau; K) dK + \int_0^{S_t} \frac{2(1 + \ln[\frac{S_t}{K}])}{K^2} P(t, \tau; K) dK \tag{47}$$

and the price of the cubic and the quartic contracts are

$$W(t, \tau) = \int_{S_t}^{\infty} \frac{6 \ln[\frac{K}{S_t}] - 3(\ln[\frac{K}{S_t}])^2}{K^2} C(t, \tau; K) dK - \int_0^{S_t} \frac{6 \ln[\frac{S_t}{K}] + 3(\ln[\frac{S_t}{K}])^2}{K^2} P(t, \tau; K) dK, \tag{48}$$

$$X(t, \tau) = \int_{S_t}^{\infty} \frac{12(\ln[\frac{K}{S_t}])^2 - 4(\ln[\frac{K}{S_t}])^3}{K^2} C(t, \tau; K) dK + \int_0^{S_t} \frac{12(\ln[\frac{S_t}{K}])^2 + 4(\ln[\frac{S_t}{K}])^3}{K^2} P(t, \tau; K) dK. \tag{49}$$

Each security price can be formulated through a portfolio of options indexed by their strikes. In addition, $\mu(t, \tau) \approx e^{r\tau} - 1 - \frac{e^{r\tau}}{2} V(t, \tau) - \frac{e^{r\tau}}{6} W(t, \tau) - \frac{e^{r\tau}}{24} X(t, \tau)$. ■

Appendix B: Simulation Experiment

To implement the simulation experiment, the stock return and volatility process are discretized as (h is some small interval):

$$S_{t+h} = S_t + \mu S_t h + \sigma_t \epsilon_t^1 \sqrt{h}, \tag{50}$$

$$\sigma_{t+h}^2 = \sigma_t^2 + \kappa(\theta - \sigma_t^2)h + v \sigma_t \epsilon_t^2 \sqrt{h}. \tag{51}$$

Simulate a time series of two independent, standard normal processes: $(\bar{\epsilon}_t^1, \bar{\epsilon}_t^2)'$, where $t = 1, 2, \dots, T$. Define

$$\mathcal{H} = \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix},$$

and generate a new vector:

$$\begin{pmatrix} \epsilon_t^1 \\ \epsilon_t^2 \end{pmatrix} = \mathcal{H}^{\frac{1}{2}} \begin{pmatrix} \bar{\epsilon}_t^1 \\ \bar{\epsilon}_t^2 \end{pmatrix}.$$

The transformed vector is a bivariate normal process with zero mean and a variance-covariance matrix of \mathcal{H} , where ϵ_t^1 and ϵ_t^2 have a correlation of ρ . Construct the time series of S_t and σ_t , $t = 1, 2, \dots, T$, based on Equations (50) and (51) and using the simulated ϵ_t^1 and ϵ_t^2 .

The initial stock price is set to be 100, and the initial value of volatility is chosen to be 10%. We initially assume that $\kappa = 2.0$, $\theta = 0.01$, $\nu = 0.1$, and $\rho = -0.50$.

For the calculations involving delta-hedged gains, the *risk-neutralized* variance process is

$$\sigma_{t+h}^2 = \sigma_t^2 + \kappa^* (\theta^* - \sigma_t^2) h + \nu \sigma_t \epsilon_t^2 \sqrt{h}, \quad (52)$$

where $\lambda_t[\sigma_t] = \lambda\sigma_t^2$, so that κ^* and θ^* are related to the physical parameters of the variance process by the relations $\kappa^* = \kappa + \lambda$, and $\theta^* = \kappa\theta/(\kappa + \lambda)$.

Suppose we set $\tau = 0.20$, then each path corresponds to 73 observations of $\{S_t, \sigma_t\}$. The delta-hedged gains, $\pi_{t,t+\tau}$, over the period, τ , is calculated using Equation (6), for a call of strike 100 and initial maturity of 0.2 years. The call price is computed as $S_t \{ \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \text{Re}[\frac{e^{-iu \ln(K)} \times f_1(u)}{iu}] du \} - K e^{-r\tau} \{ \frac{1}{2} + \frac{1}{\pi} \int_0^\infty \text{Re}[\frac{e^{-iu \ln(K)} \times f_2(u)}{iu}] du \}$, where the characteristic functions, f_1 and f_2 , are displayed in Heston [1993, Equation (17)]. For simplicity, the interest rate and the dividend yield are assumed to be zero. ■

References

- Anderson, T., L. Benzoni, and J. Lund, 2002, "An Empirical Investigation of Continuous-Time Equity Return Models," *Journal of Finance*, 57, 1239–1284.
- Bakshi, G., C. Cao, and Z. Chen, 1997, "Empirical Performance of Alternative Option Pricing Models," *Journal of Finance*, 52, 2003–2049.
- Bakshi, G., C. Cao, and Z. Chen, 2000, "How Often Does the Call Move in the Opposite Direction to the Underlying?," *Review of Financial Studies*, 13, 549–584.
- Bakshi, G., N. Kapadia, and D. Madan, 2003, "Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options," *Review of Financial Studies*, 16, 101–143.
- Bates, D., 2000, "Post-87 Crash Fears in S&P 500 Futures Options," *Journal of Econometrics*, 94, 181–238.
- Benzoni, L., 1999, "Pricing Options Under Stochastic Volatility: An Econometric Analysis," mimeo, University of Minnesota.
- Bertsimas, D., L. Kogan, and A. Lo, 2000, "When is Time Continuous," *Journal of Financial Economics*, 55, 173–204.
- Buraschi, A., and J. Jackwerth, 2001, "The Price of a Smile: Hedging and Spanning in Option Markets," *Review of Financial Studies*, 14, 495–527.
- Chernov, M., R. Gallant, E. Ghysels, and G. Tauchen, 2002, "Alternative Models for Stock Price Dynamics," forthcoming in *Journal of Econometrics*.
- Chernov, M., and E. 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.
- Christensen, B., and N. Prabhala, 1998, "The Relation Between Implied and Realized Volatility," *Journal of Financial Economics*, 50, 125–150.
- Coval, J., and T. Shumway, 2001, "Expected Option Returns," *Journal of Finance*, 56, 983–1009.
- Duan, J. C., I. Popova, and P. Ritchken, 2002, "Option Pricing Under Regime-Switching," *Quantitative Finance*, 2.

- Dumas, B., J. Fleming, and R. Whaley, 1998, "Implied Volatility Functions: Empirical Tests," *Journal of Finance*, 53, 2059–2106.
- Eraker, B., M. Johannes, and N. Polson, 2003, "The Impact of Jumps in Volatility and Returns," forthcoming in *Journal of Finance*.
- Figlewski, S., 1989, "Options Arbitrage in Imperfect Markets," *Journal of Finance*, 44, 1289–1311.
- French, K., W. Schwert, and R. Stambaugh, 1987, "Expected Stock Returns and Volatility," *Journal of Financial Economics*, 19, 3–29.
- Glosten, L., R. Jagannathan, and D. Runkle, 1993, "On the Relation Between the Expected Value and the Volatility of the Nominal Excess Returns on Stock," *Journal of Finance*, 48, 1779–1801.
- Greene, W., 1997, *Econometric Analysis*, 3rd edition, MacMillan, New York.
- Hansen, L., 1982, "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 50, 1029–1084.
- Heston, S., 1993, "A Closed Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options," *Review of Financial Studies*, 6, 327–343.
- Heston, S., and S. Nandi, 2000, "A Closed-Form GARCH Option Valuation Model," *Review of Financial Studies*, 13, 585–625.
- Hull, J., and A. White, 1987, "The Pricing of Options with Stochastic Volatilities," *Journal of Finance*, 42, 281–300.
- Jackwerth, J., 2000, "Recovering Risk Aversion from Options Prices and Realized Returns," *Review of Financial Studies*, 13, 433–451.
- Jackwerth, J., and M. Rubinstein, 1996, "Recovering Probability Distributions from Option Prices," *Journal of Finance*, 51, 1611–1631.
- Jones, C., 2002, "The Dynamics of Stochastic Volatility: Evidence from Options and the Underlying Market," forthcoming in *Journal of Econometrics*.
- Merton, R., 1973, "Theory of Rational Option Pricing," *Bell Journal of Economics and Management Science*, 4, 41–83.
- Merton, R., 1976, "Option Pricing When Underlying Stock Returns are Discontinuous," *Journal of Financial Economics*, 3, 125–144.
- Milstein, G., 1995, *Numerical Integration of Stochastic Differential Equations*, Kluwer Academic, Boston.
- Pan, J., 2002, "The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study," *Journal of Financial Economics*, 63, 3–50.
- Poteshman, A., 1998, "Estimating a General Stochastic Variance Model from Option Prices," mimeo, University of Illinois.
- Rubinstein, M., 1994, "Implied Binomial Trees," *Journal of Finance*, 49, 771–818.
- Stein, J., 1989, "Overreactions in the Options Market," *Journal of Finance*, 44, 1011–1023.