Volatility Risk Premiums Embedded in Individual Equity Options: *Some New Insights*

GURDIP BAKSHI AND NIKUNJ KAPADIA

GURDIP BAKSHI
is Smith Professor of Finance at the Smith School of Business of the University of Maryland, in College Park, MD.
gbakshi@rhsmith.umd.edu

NIKUNJ KAPADIA
is an associate professor of finance at the Isenberg School of Management of the University of Massachusetts at Amherst, MA.
nkapadia@som.umass.edu

The research indicates that index option prices incorporate a negative volatility risk premium, thus providing a possible explanation of why Black-Scholes implied volatilities of index options on average exceed realized volatilities. This examination of the empirical implication of a market volatility risk premium on 25 individual equity options provides some new insights.

While the Black-Scholes implied volatilities from individual equity options are also greater on average than historical return volatilities, the difference between them is much smaller than for the market index. Like index options, individual equity option prices embed a negative market volatility risk premium, although much smaller than for the index option—and idiosyncratic volatility does not appear to be priced.

These empirical results provide a potential explanation of why buyers of individual equity options leave less money on the table than buyers of index options.

In Bakshi and Kapadia [2003] we show there is a negative volatility risk premium in index options. One crucial impact is to make index options more expensive. A negative market volatility risk premium provides at least a partial explanation for the finding that index implied volatilities are typically greater than realized volatilities (Jackwerth and Rubinstein [1996]). In essence, adding options to a market portfolio will help hedge market risks as markets tend to become more volatile when the stock market falls, consistent with a negative volatility risk premium.

We investigate the pricing of market volatility risk in individual equity options. There are several reasons why it is important to extend our evidence on the index options market to individual equity options. First, given that stock returns have a significant market component, the presence of a market volatility risk premium has implications for how individual equity options are priced. We may be able to verify that market volatility risk is compensated and that results for the index option market are not driven by other factors (say, demand for index options for portfolio insurance purposes).

Second, it is of economic importance to understand the extent of the volatility risk premium embedded in individual stock options. Given that individual risk-neutral distributions are systematically different from the market index, how volatility risk is priced in individual options can give us additional insights into the pricing structure of individual equity options (see Bakshi, Kapadia, and Madan [2003]).

Following the theoretical arguments in Bakshi and Kapadia [2003], we consider the gains on a delta-hedged option portfolio—a portfolio of a long call position hedged by a short position in the stock, with the net investment earning the risk-free interest rate. If volatility risk is not priced, average delta-hedged gains are zero even when volatility is stochastic. If volatility risk is priced, then the sign and the magnitude of the average delta-hedged gains are determined by the volatility risk premium. A testable implication of a non-zero market volatility risk premium is that delta-hedged gains are correlated with the level of market volatility and not with idiosyncratic return volatility.

We apply our framework to an empirical examination of the pricing of market volatility risk in 25 individual equity options. Our first finding is that, on average, near-money Black-Scholes implied volatilities from individual equity options are greater than historical realized volatilities. More important, there is a much smaller difference between implied and realized volatilities for individual equity options than for the index.

Second, delta-hedged gains of individual equity options are more negative than positive. Almost twice as many firms have significantly negative gains as significantly positive gains. Over all firms, on average, the delta-hedging strategy loses a statistically significant 0.03% of underlying asset value. The same delta-hedging strategy for the index loses 0.07% of the underlying index level.

Third, individual equity delta-hedged gains are significantly negatively correlated with the level of market volatility. Moreover, the market volatility subsumes the effect of the firm's own volatility, and idiosyncratic volatility does not appear to be priced. Our results are consistent with the implication of a non-zero volatility risk premium, particularly a negative market volatility risk premium, although much smaller than for index options.

What are the economic implications of these results? The difference between Black–Scholes implied volatility and realized volatilities indicates that, like the buyers of index options, buyers of individual equity options also lose money and, in other words, leave money on the table. The analysis of delta-hedged gains provides a reason why—individual equity options, like index options, also incorporate a negative market volatility risk premium. Our estimates indicate the volatility risk premium is much lower for individual equities, thus providing an explanation of why implied and realized volatility are empirically closer in individual option markets.

I. BASIC MODEL

The theoretical framework is a variant of that adopted in Bakshi and Kapadia [2003]. We denote the stock price and the variance of firm *i* as $S_i(t)$, and $V_i(t)$, respectively, and the market index variance as $V_m(t)$. Under the physical probability measure, assume that:

$$\frac{dS_i(t)}{S_i(t)} = \mu_i[S_i, V_i]\,dt + \sqrt{V_i(t)}\,dW^1(t) \qquad (1)$$

$$V_i(t) = \beta_i V_m(t) + Z_i(t), \qquad \beta_i > 0 \qquad (2)$$

and

$$dV_m(t) = \theta[V_m]\,dt + \eta[V_m]\,dW^2(t) \qquad (3)$$

To explain the stock price dynamics assumptions outlined in (1)-(3), we first note that Equation (2) implies a single-factor model of individual stock variance. Equation (2) is based on the assumption that stock returns have a market component and an idiosyncratic component, as in $R_i(t) = \hat{\alpha}_i + \hat{\beta}_i R_m(t) + \hat{\varepsilon}_i(t)$. Specifically, if the idiosyncratic stock variance $Z_i(t)$ is uncorrelated with market index variance, then $\beta_i \equiv \hat{\beta}_i^2$ is the sensitivity of individual variance with respect to the variance of the market index. That is, we are essentially staying within the Black-Scholes framework, with a modification allowing for stochastic volatility in the individual stock price process in Equation (1).

See Bakshi and Kapadia [2003] for a framework that relates the losses on delta-hedged portfolios to return jumps. Given the low negative risk-neutral skewness found in individual equity options, Merton [1976] type return jumps are omitted to maintain focus on the volatility risk premium.

Equation (3) generically specifies the market variance as a one-dimensional diffusion with drift and diffusion coefficients given by $\theta[V_m]$ and $\eta[V_m]$. While certain choices for $\theta[V_m]$ and $\eta[V_m]$ can lead to empirically unappealing variance dynamics (i.e., arithmetic and non-mean-reverting), we nonetheless keep the functional form of $\theta[V_m]$ and $\eta[V_m]$ unspecified, with the understanding that suitable $\theta[V_m]$ and $\eta[V_m]$ imply a well-specified variance process.

If $\theta[V_m] \equiv \theta - \kappa\, V_m$ and $\eta[V_m] \equiv \eta\sqrt{V_m(t)}$, Equation (3) admits mean reversion in market return variance and a stationary distribution for both $V_m(t)$ and $V_i(t)$ (see Heston [1993]). Let ρ_i be the correlation between the standard Brownian motions $W^1(t)$ and $W^2(t)$.

For ease of analysis, assume that idiosyncratic return variance, $Z_i(t)$, is a constant for all t so that we can set $dZ_i(t) = 0$ in $dV_i(t) = \beta dV_m(t) + dZ_i(t)$. With constant Z, individual return variance obeys a one-dimensional diffusion. As we show below, however, our key result will be unaffected with stochastic $Z_i(t)$, provided $Z_i(t)$ is unpriced (i.e., uncorrelated with the pricing kernel). (In a later empirical exercise we reject that idiosyncratic return volatility is priced.)

Denote $C_i(t; K, \tau)$ as the call option on the individual stock with strike K and maturity τ. Then by Itô's lemma:

$$C_i(t+\tau) = C_i(t) + \int_t^{t+\tau} \Delta_i(u)\, dS_i(u) + \int_t^{t+\tau} \frac{\partial C_i}{\partial V_i}\, dV_i(u) + \int_t^{t+\tau} b_i(u)\, du, \tag{4}$$

where

$$b_i(u) \equiv \frac{\partial C_i}{\partial u} + \frac{1}{2} V_i\, S_i^2 \frac{\partial^2 C_i}{\partial S_i^2} + \frac{1}{2}\beta_i^2\, \eta[V_m]^2 \frac{\partial^2 C_i}{\partial V_i^2} + \rho\, \beta_i\, \eta[V_m]\, \sqrt{V_i}\, S_i \frac{\partial^2 C_i}{\partial S_i \partial V_m} \tag{5}$$

and

$$\Delta_i \equiv \partial C_i / \partial S_i \tag{6}$$

is the call option delta. In the empirical implementation, we treat the Black-Scholes delta as a close enough approximation to the true delta that it can be used as a proxy.

The valuation equation that determines the price of the call option is:

$$\frac{1}{2} V_i\, S_i^2 \frac{\partial^2 C_i}{\partial S_i^2} + \frac{1}{2}\eta[V_m]^2\, \beta_i^2 \frac{\partial^2 C_i}{\partial V_i^2} + \rho\, \beta_i\, \eta[V_m]\, \sqrt{V_i}\, S_i \frac{\partial^2 C_i}{\partial S_i \partial V_i} + r S_i \frac{\partial C_i}{\partial S_i} + \beta_i\,(\theta[V_m] - \lambda[V_m]) \frac{\partial C_i}{\partial V_i} + \frac{\partial C_i}{\partial t} - r\, C_i = 0 \tag{7}$$

where

$$\lambda\,[V_m] \equiv -\mathrm{Cov}_t \left(\frac{dm(t)}{m(t)}, dV_m(t) \right) \tag{8}$$

represents the price of volatility risk for a pricing kernel process m_t and $\mathrm{Cov}_t(,.)$ is a conditional covariance operator divided by dt.

Note that if volatility is non-stochastic, as in the basic Black-Scholes model, volatility risk is zero, and it does not matter whether volatility is high or low, as the resulting delta hedge is riskless in theory. Allowing volatility to be stochastic exposes the investor to random variability in volatility through its covariance with changes in the pricing kernel, as made precise in Equation (8).

If the market volatility risk premium is non-zero, as empirically shown by Bakshi and Kapadia [2003], and $\beta_i > 0$, the individual volatility risk premium will have the same sign as the market volatility risk premium. This characterization would hold if we define volatility risk as:

$$-\mathrm{Cov}_t \left(\tfrac{dm(t)}{m(t)}, d\sqrt{V_m(t)} \right)$$

which, by Itô's lemma, is the same as:

$$-\tfrac{1}{2\sqrt{V_m}}\, \mathrm{Cov}_t \left(\tfrac{dm(t)}{m(t)}, dV_m(t) \right)$$

In particular, if $\lambda[V_m]$ is assumed proportional in V_m, then $\lambda[\sqrt{V_m}]$ is proportional to $\sqrt{V_m}$. Thus, in what follows, we adopt the convention that volatility risk is as specified in (8), and that volatility and variance are interchangeable in the discussion.

Combining (7) and substituting $b_i(u)$ in the stochastic differential Equation (4), we obtain

$$C_i(t+\tau) - C_i(t) = \int_t^{t+\tau} \frac{\partial C_i}{\partial S_i}\, dS_i + \int_t^{t+\tau} r \left(C_i - \frac{\partial C_i}{\partial S_i} S_i \right) du + \int_t^{t+\tau} \beta_i\, \lambda[V_m] \frac{\partial C_i}{\partial V_i} du + \int_t^{t+\tau} \beta_i\, \eta[V_m] \frac{\partial C_i}{\partial V_i}\, dW^2 \tag{9}$$

Define the delta-hedged gains, $(t,\ t + \tau)$, as the gain or loss on a delta-hedged option position (where the net investment earns the risk-free rate):

$$\Pi_i(t, t+\tau) \equiv C_i(t+\tau) - C_i(t) - \int_t^{t+\tau} \Delta_i\, dS_i - \int_t^{t+\tau} r\,(C_i - \Delta_i\, S_i)\; du \tag{10}$$

Then, from (9) and (10), we can write the delta-hedged gains as:

$$\Pi_i(t, t+\tau) = \int_t^{t+\tau} \beta_i\, \lambda[V_m] \frac{\partial C_i}{\partial V_i}\, du + \int_t^{t+\tau} \beta_i\, \eta[V_m] \frac{\partial C_i}{\partial V_i}\, dW^2 \tag{11}$$

and from the martingale property of the Itô integral:

$$\mathrm{E}_t(\Pi_i(t,t+\tau)) = \int_t^{t+\tau} \mathrm{E}_t\left(\beta_i \times \lambda[V_m] \times \frac{\partial C_i}{\partial V_i}\right) du \quad (12)$$

where $\mathrm{E}_t(\cdot)$ is the expectations operator under the physical probability measure.

The testable implication of (12) is that the magnitude and the sign of the individual equity delta-hedged gains are related to the sign and magnitude of $\lambda[V_m]$. As $\lambda[.]$ is some function of market variance, V_m, it follows that, if market volatility risk is priced, the delta-hedged gains will be related to the level of market variance. In other words, we can deduce the impact of the market volatility risk premium by the relation between individual equity delta-hedged gains and market variance.

Equation (12) also holds if we relax the assumption that idiosyncratic return volatility is constant. To allow for stochastic idiosyncratic volatility, let:

$$dZ_i(t) = \theta_i^z[Z_i]\, dt + \eta_i^z[Z_i]\, dW^3(t) \quad (13)$$

Then under the assumption that standard Brownian motion $W^3(t)$is independent of all sources of stochastic variation, we can derive, as before, the delta-hedged gains as:

$$\Pi_i(t,t+\tau) = \int_t^{t+\tau} \beta_i\, \lambda[V_m]\, \frac{\partial C_i}{\partial V_i}\, du +$$

$$\int_t^{t+\tau} \beta_i\, \eta[V_m]\, \frac{\partial C_i}{\partial V_i}\, dW^2 + \int_t^{t+\tau} \eta_i^z[Z_i]\, \frac{\partial C_i}{\partial V_i}\, dW^3 \quad (14)$$

Because idiosyncratic volatility is not priced, Equation (12) still applies; stochastic idiosyncratic volatility simply adds noise to the delta-hedged gains. Moreover, this fact provides an additional test of whether market volatility risk is priced. Mean delta-hedged gains must be correlated with market volatility, and not with firm-specific idiosyncratic volatility. Our analysis can be extended in a straightforward manner to allow for multifactor models of individual return volatility.

In general, the relation between the delta-hedged gains and market volatility in (12) may be of any functional form. For the square root process of Heston [1993], however, this relation is strikingly simple. As before, let $\theta[V_m] \equiv \theta - \kappa V_m$, $\eta[V_m] \equiv \eta\sqrt{V_m(t)}$, and $\lambda[V_m] \equiv \lambda V_m(t)$. Then we show in Bakshi and Kapadia [2003] that, for at-the-money options, the *scaled* delta-hedged gains, defined as $\Pi_i(t, t+\tau)(S_i(t)$, are related linearly to the level of the market volatility. The scaling of the delta-hedged gains by the price of the underlying asset ensures that we can compare delta-hedged gains over time.

The linearity of the scaled delta-hedged gain implies that a linear regression can be used to test the implications of the market volatility risk premium. Thus, there is a relatively straightforward test that allows us to answer the two questions of importance: Is market volatility risk priced, and, if so, what is the sign and strength of the risk premium?

II. INDIVIDUAL EQUITY OPTIONS

Our empirical tests use bid-ask call option quotes on 25 individual stocks and the S&P 500 index obtained from the Berkeley Options Database. These options are traded on the Chicago Board Options Exchange and have American-style exercise. For each of the 1,258 days in the sample period of January 1, 1991–December 31, 1995, we retain the last quote prior to 3:00 PM (CST). The individual stock options are identified by ticker symbol and name in Exhibit 1. This sample includes the largest stocks, as their options are likely to be more liquid.

Several filters are employed to construct the call sample. First, we screen the data to eliminate bid-ask option pairs with missing quotes or zero bids. Second, we remove option prices violating arbitrage restrictions, $C(t,\tau;K) < S(t)$ or $C(t,\tau;K) > S(t)$ – PVD[D] – PVD[K] present value function PVD[.] and dividends D. Third, we eliminate options with fewer than 14 or more than 30 days remaining to expiration. Finally, we use close-to-the-money options within a moneyness range of $-2.50\% \leq y(t) - 1 \leq 2.50\%$, where $y(t) \equiv S(t)e^{r\tau}/K$ is option moneyness.

Firm-specific dividends are obtained from the Center for Research in Security Prices and are assumed known. The source of S&P 500 dividends is the S&P Information Bulletin. Following convention, the current stock price is adjusted by subtracting the present value of dividends. As in Bakshi and Kapadia [2003], we interpolate the interest rate to match option maturity using overnight and three-month Eurodollar interest rates.

For our calculations involving realized volatility, we use a measure of sample standard deviation, which is computed as:

$$\mathrm{VOL}_{t-\tau,t} = \sqrt{\frac{252}{\tau} \sum_{n=t-\tau}^{t} R_{n-1,n}^2} \quad (15)$$

where τ is set to 30/360 for monthly volatility. We do not subtract the sample mean return as this estimate of expected return can be unreasonable.

As empirical results using GARCH estimation are essentially similar, we do not report them (see the robust-

EXHIBIT 1

Implied Volatility versus Realized Volatility—1991–1995

Ticker	Firm	Panel A: All Options				Panel B: Omitting Dividends			
		OBS	IVOL	$VOL_{t,t+\tau}$	$VOL_{t-30,t}$	OBS	IVOL	$VOL_{t,t+\tau}$	$VOL_{t-30,t}$
AIG	American Int'l	668	21.5	18.6	19.6	513	21.4	18.9	19.2
AIT	Ameritech	431	18.4	18.3	18.3	379	18.2	18.0	18.5
AN	Amoco	354	17.8	17.8	18.2	231	17.2	18.0	17.8
AXP	American Exp.	323	27.8	25.8	26.2	237	27.8	27.0	26.8
BA	Boeing	353	23.8	21.5	21.7	238	23.0	21.6	21.3
BAC	BankAmerica	307	26.9	24.1	24.9	224	26.1	24.1	23.8
BMY	Bristol Myers	452	19.7	17.8	19.0	349	19.6	17.6	19.2
CCI	Citigroup	244	29.3	28.8	30.9	215	29.2	29.4	31.0
DD	Du Pont	342	22.7	21.3	20.6	215	21.7	20.0	19.5
DIS	Walt Disney	499	26.2	23.1	24.1	327	26.5	22.8	25.4
F	Ford Motor	284	28.2	27.7	28.3	261	27.9	27.9	28.2
GE	General Electric	514	18.8	17.4	17.3	420	18.5	17.4	17.1
GM	General Motors	279	28.7	28.3	27.9	170	28.1	27.6	27.7
IBM	Int. Bus. Mach.	529	25.9	23.4	24.0	365	25.4	24.1	23.4
JNJ	J & J	476	22.9	22.1	22.6	315	22.6	21.7	21.9
KO	Coca Cola	374	22.0	20.4	20.4	279	22.3	20.9	20.5
MCD	MacDonald's	329	21.5	20.7	21.6	278	21.4	21.0	21.3
MMM	Minn. Mining	603	19.4	17.2	17.4	451	18.5	17.3	17.1
MOB	Mobil	546	17.5	16.7	16.7	407	17.1	16.8	17.0
MRK	Merck	483	23.9	21.5	22.1	365	23.9	22.5	21.9
PEP	Pepsico	261	21.5	20.6	21.7	172	21.6	21.5	21.1
SLB	Schlumberger	396	23.7	23.3	23.9	317	23.6	23.4	23.8
T	AT&T	298	18.5	17.9	18.5	255	18.3	17.7	18.4
WMT	Walmart	276	25.3	22.9	23.3	208	25.5	23.6	23.1
XRX	Xerox	572	24.3	21.3	22.3	440	24.2	22.0	21.6
SPX	S&P 500	2990	12.8	9.5	9.5	-	-	-	-

The table reports (i) Black-Scholes implied volatility for near-money individual equity calls (denoted IVOL), (ii) realized volatility (denoted $VOL_{t,t+\tau}$) and (iii) prior 30-day realized volatility (denoted $VOL_{t-30,\tau}$). The implied volatility is computed by equating the market option price to the Black-Scholes model price. The sample period is January 1991 to December 1995. In the implied volatility calculation, dividends are assumed known, discounted and subtracted from the stock price. All options have remaining days to maturity of 15-30 days. Near-money calls are defined to have a moneyness between $-2.50\% \leq y_i(t) - 1 \leq 2.50\%$ where $y_i(t) \equiv S_i(t)e^{r\tau}/K_i$ and K_i is the strike of the option. Panel A reports the results using all call options; Panel B reports the results omitting options if the stock pays a cash dividend. OBS is the number of observations.

ness exercises in Bakshi and Kapadia [2003]). The measure of standard deviation in (15) is convenient, as the rolling procedure produces estimates whose estimation error is serially uncorrelated through time for non-overlapping periods.

III. IMPLIED AND REALIZED VOLATILITIES

To examine whether volatility risk is priced in individual equities, we first investigate the relation between implied and realized volatilities. Jackwerth and Rubinstein [1996] document that Black-Scholes implied volatilities for at-the-money index options are on average greater than realized volatilities. This empirical finding has proven difficult to reconcile by just relaxing the constant-volatility assumption, by allowing, for example, stochastic volatility or jumps in the stock return dynamics (see, for example, Bakshi, Cao, and Chen [1997] and Bates [2000]). Bakshi and Kapadia [2003] and Buraschi and Jackwerth [2001] suggest one possible answer—that market volatility risk is priced in equity options.

Specifically, in Bakshi and Kapadia [2003] we reason that option prices incorporate a negative volatility risk premium. The empirical tests are motivated by the finding that shocks to market volatility are negatively correlated with market returns (e.g., French, Schwert, and Stambaugh [1987]). The negative correlation implies that market volatility increases when the market return is negative. In this case, including options in a portfolio will help hedge market risk, as the option vega is positive. The hedging motive makes investors willing to pay a risk premium for a long option position, implying a negative volatility risk premium.

A negative volatility risk premium increases the option price, resulting in an implied volatility that is higher than expected future volatility. More precisely, the drift of the risk-neutral volatility process will exceed the drift under the physical probability measure. Because individual volatilities are generally positively correlated with market volatility

EXHIBIT 2
Delta-Hedged Gains—1991–1995

			Panel A		Panel B		
			Magnitude of $\Pi(t,t+\tau)/S(t)$		**Magnitude of** $\Pi(t,t+\tau)/C(t)$		
Ticker	**Firm**	**OBS**	Average	Median	Average	Median	$1_{\Pi<0}$ (%)
AIG	American Int'l	565	-0.20	-0.30	-9.29	-14.99	67.08
AIT	Ameritech	409	0.04	-0.11	3.21	-6.33	57.70
AN	Amoco	256	0.12	0.06	8.65	3.71	44.14
AXP	American Exp.	255	0.05	0.01	0.78	0.41	48.63
BA	Boeing	255	-0.04	-0.16	1.66	-7.97	57.65
BAC	BankAmerica	242	-0.17	-0.17	-11.01	-9.21	64.46
BMY	Bristol Myers	383	-0.12	-0.23	-5.50	-12.23	67.10
CCI	Citigroup	235	0.03	-0.06	2.27	-2.65	51.49
DD	Du Pont	228	-0.14	-0.21	-5.17	-10.06	63.16
DIS	Walt Disney	358	-0.26	-0.36	-9.48	-13.03	68.99
F	Ford Motor	286	0.22	0.07	10.44	2.49	46.85
GE	General Electric	460	-0.05	-0.08	-1.20	-3.68	54.57
GM	General Motors	185	-0.08	-0.15	-1.30	-5.36	56.22
IBM	Int. Bus. Mach.	396	-0.05	-0.19	-0.59	-8.92	63.64
JNJ	J & J	342	0.06	-0.09	1.27	-4.05	55.26
KO	Coca Cola	304	-0.11	-0.11	-1.31	-4.46	59.87
MCD	MacDonald's	303	0.10	0.06	4.58	3.87	44.88
MMM	Minn. Mining	488	0.05	-0.11	4.65	-7.28	60.25
MOB	Mobil	446	0.01	-0.00	0.44	-0.15	50.45
MRK	Merck	397	-0.13	-0.22	-4.26	-9.15	66.50
PEP	Pepsico	187	0.11	-0.02	5.62	-0.78	50.80
SLB	Schlumberger	345	0.11	0.00	11.76	0.12	49.57
T	AT&T	281	-0.01	-0.06	-2.16	-4.20	55.52
WMT	Walmart	224	-0.05	-0.21	-0.16	-9.90	59.38
XRX	Xerox	490	-0.05	-0.21	-2.68	-10.96	61.02
SPX	S&P 500	2990	-0.07	-0.10	-3.31	-5.98	65.79

Panel A: Delta-hedged gains normalized by the stock; Panel B: Delta-hedged gains normalized by the option price. $1_{\Pi<0}$ *is the proportion of* Π *that is less than zero. Individual equity calls are excluded if the underlying stock pays a dividend within the maturity of the option.*

[see Equation (2)], we might expect the individual implied volatilities to similarly exceed realized volatilities.

Panel A of Exhibit 1 reports the Black-Scholes implied volatilities and the realized volatilities for all options. Consistent with the assumption that volatility risk is priced, the implied volatilities of individual equity options tend to be higher than the realized volatilities. For example, the average annualized implied volatility for GE is 18.8% as compared with an annualized historical volatility of 17.4% realized over the remaining lifetime of the option. This conclusion is robust even if return volatility is measured using returns over the past 30 calendar days.

For every option in the sample, however, the difference between implied and historical volatilities is *less* than that for S&P 500 index options. For instance, the average difference between implied and realized volatility for SPX calls is 3.3% (on an annualized basis), while the average across the 25 stocks in our sample is only 1.5%. Given that the options on individual stocks are priced at a higher level of volatility, this difference between the implied and the realized volatility has a considerably smaller price impact for individual equity options (than for the market index).

One possible concern with this result is that individual equity options are American-style while the SPX option is European. To assess the impact of early exercise on our results, we can use the fact that if there are no dividends paid in the remaining maturity of the call, the early exercise of the American call will not be optimal. Guided by this result, we also compare implied volatilities of options that have a dividend in the remaining time to maturity to the implied volatility when there is no dividend prior to maturity. This exercise shows that the early exercise premium is equivalent to about 2 percentage points of volatility for our sample of short-term near-money calls.

To eliminate the impact of early exercise on our results in the empirical work that follows, we eliminate all call observations where the stock pays a cash dividend over the remaining lifetime of the contract. Panel B of Exhibit 1 reports the results of the option sample omitting dividends. Because of quarterly dividend payouts, about 25% of the original individual equity option sample is eliminated.

The average difference between the Black-Scholes implied and realized volatility across the 25 firms is now 1.07%, lower than the estimate in Panel A. Note that for

23 of 25 firms, the average implied volatility still exceeds realized volatility. Moreover, this finding is also robust when the return volatility is estimated using the prior 30 days' stock returns.

Overall, two chief conclusions can be drawn from Exhibit 1. First, the pricing of individual equity options and index options is consistent with the notion that implied volatilities are, on average, higher than realized volatilities. This suggests that volatility risk is negatively priced in both individual and index option markets.

Second, the difference between implied and realized volatilities is far more pronounced for index options than for individual equity options. It appears that individual equity option buyers in general leave money on the table, but leave more for index options.

IV. INSIGHTS FROM INDIVIDUAL EQUITY OPTIONS

Can a market volatility risk premium help explain these findings? To examine the implications of a priced market volatility risk factor, we first construct delta-hedged gains for every option in our sample. The option position consists of a long call bought at time t and hedged discretely until expiration, $t + \tau$. The total delta-hedged gain for each option to the maturity date is calculated as:

$$\begin{aligned}\Pi(t, t+\tau) &= \max(S(t+\tau) - K, 0) - \\ & C(t) - \sum_{n=0}^{N-1} \Delta(t_n)(S(t_{n+1}) - S(t_n)) - \\ & \sum_{n=0}^{N-1} r\,(C(t) - \Delta(t_n) S(t_n))\,\frac{\tau}{N} \qquad (16)\end{aligned}$$

where $t_0 = t,\ t_N = t + \tau$ is the maturity date, and $\Delta(t_n)$ is the hedge ratio at t_n (recomputed daily). For tractability, Δ_{tn} is computed as the Black-Scholes hedge ratio, $\Delta_{tn} = \mathcal{N}\ [d_1(S_{tn},\ _{tn})]$ where $\mathcal{N}[\ .\ ]$is the cumulative normal distribution, and

$$d_1 \equiv \frac{1}{\mathrm{VOL}_{t,t+\tau}\sqrt{\tau_n}} \log(y_n) + \frac{1}{2}\mathrm{VOL}_{t,t+\tau}\sqrt{\tau_n}.$$

Bakshi and Kapadia [2003] note that use of the Black-Scholes hedge ratio as an approximation of the true hedge ratio does not significantly alter the conclusions, and the results are also robust to alternative estimates of return volatility in $\Delta(t_n)$.

Exhibit 2 reports the magnitude of delta-hedged gains for 25 individual stocks and the SPX, again excluding equity option observations on dates the firm paid a dividend during the remaining lifetime of the option. To make the delta-hedged gains comparable across the time series and the cross-section, we express the normalized delta-hedged gains as $\Pi_i(t, \tau)/S_t$ *and* $\Pi_i(t, \tau)/C_t$.

The delta-hedging strategy for the SPX loses money. On average, SPX calls lose 0.07% of the value of the index. In terms of the value of the option, the average loss is 3.31%. This loss is both statistically and economically significant; given the traded volume of index options, the dollar loss amounts to several $100 million over the time period.

Why should buyers of options be willing to leave money on the table? As we have emphasized, negative delta-hedged gains are consistent with a negative market volatility risk premium $\lambda[V_m] < 0$. A negative volatility risk premium increases the option price in comparison with its price when $\lambda[V_m] \equiv 0$. Because of the negative correlation between market index returns and market index volatility, buyers of options may be willing to pay a premium because a long position in volatility helps hedge marketwide risk.

It is important to realize that the negative delta-hedged gains for the index option do not necessarily translate into negative delta-hedged gains for the individual equity option. Although we would expect a negative market volatility risk premium to impart a negative bias to the delta-hedged gains for individual equity options, market volatility is merely one component of the firm's total volatility. The impact of the market volatility risk premium on the individual firm would depend on both the relation between the firm's total volatility and the market volatility and whether non-market components of volatility are priced. The importance of market volatility in determining the distribution of delta-hedged gains for individual equity options can be determined only empirically, and is the focus of our interest.

Exhibit 2 indicates that the majority of the stocks have negative delta-hedged gains. The average delta-hedged gain is negative for 14 of the 25 individual equity options, and 7 of the firms have gains of significantly less than zero at the 99% level (standard errors are not reported, but are available upon request). Only four firms have gains that are significantly positive. Our results remain robust when the median is used as a measure of central tendency.

To eliminate biases that may be caused by a few outliers, we also examine the relative outcomes of negative and positive gains across the firms. The last column of Exhibit

E X H I B I T 3

Time Series Relationship Between Delta-Hedged Gains and Volatility—1991–1995

Dependent Variable	T	Ω_0 $\times 10^{-2}$	Ω_1	R^2 (in %)
$\overline{\text{GAINS}}$	60	0.11 [1.49]	-0.018 [-2.18]	3.51
GAINS_m	59	0.14 [1.32]	-0.030 [-2.92]	8.43

$\overline{\text{GAINS}}$: *Average delta-hedged gains for month* t *over the 25 individual stocks in the sample;* GAINS_m*: Month* t *delta-hedged gains for the SPX; t-statistics in brackets.*

2 displays the $1_{\Pi < 0}$ statistic, which measures the frequency of negative outcomes. Of the 25 individual stocks, 20 have negative delta-hedged gains more than 50% of the time. On average, over the 25 stocks in our sample, the delta-hedging strategy has a loss of 0.03% of the underlying asset value.

Although the delta-hedging trading strategy loses money for both the index and individual equity options, the loss on average for individual equity options is far less than that for the index. Overall, the evidence from individual firms is consistent with that observed for the index option in that, on average, delta-hedged gains are negative.

V. DELTA-HEDGED GAINS, MARKET VOLATILITY, AND IDIOSYNCRATIC VOLATILITY

Are the observed delta-hedged gains for individual equity options consistent with a market volatility risk premium? For negative delta-hedged gains to be consistent with a negative market volatility risk premium, they must be negatively correlated with the level of market volatility and unrelated to idiosyncratic return volatility.

An important theoretical implication of Equation (12) is that the extent of the delta-hedged gains depends on the level of market volatility. We show in Bakshi and Kapadia [2003] that, under mild assumptions, the relation between the scaled delta-hedged gains, $\Pi_i(t, t + \tau)/S_i(t)$, for near-money options is linear in $\sqrt{V_m(t)}$. Thus, we can test the relation between market volatility and delta-hedged gains using a linear regression framework.

For this test, we choose the closest-to-the-money option (within a 2.5% moneyness range) with maturity of exactly 30 days. Once again, options with dividends paid over the period are eliminated. The number of observations for each firm ranges from 19 to 40, with a total of 610 observations. The market volatility is estimated on the basis of returns realized over the previous 30 calendar days.

We investigate the implication of a market volatility risk premium via two tests. First, we average the delta-hedged gains for every month over the 25 stocks in the sample, and then regress it on market volatility (letting $\text{VOL}_m(t) = \text{VOL}_{t-30,t}$):

$$\overline{\text{GAINS}}(t) = \Omega_0 + \Omega_1 \text{VOL}_m(t) + \epsilon_t \qquad (17)$$

where $\overline{\text{GAINS}}(t) = \frac{1}{25}\sum_{i=1}^{25} \text{GAINS}_i(t)$.

As firms have different dividend payment cycles, averaging across all firms in any period allows us to construct a time series of observations that is complete over the 60 months of our data period. Although this test does not allow us to discern firm-specific components of the delta-hedged gains, it does allow us to test whether, on average, market volatility risk is priced in individual equity options.

Exhibit 3 reports the regression results for the specification in (17). We observe that the average delta-hedged gain is, in fact, negatively correlated with the prior month market volatility. The estimate of Ω_1 is –0.018 and is statistically significant.

For comparison, we also report the regression results when the dependent variable is the delta-hedged gain for the SPX option:

$$\text{GAINS}_m(t) = \Omega_0 + \Omega_1 \text{VOL}_m(t) + \epsilon_t^* \qquad (18)$$

The estimate of the slope coefficient is –0.03, again statistically significant. The negative sign of the slope coefficient in both regressions is consistent with the presence of a negative market volatility risk premium.

An alternative econometric specification could allow for both a firm-specific component and a market component in the delta-hedged gains. Recall from the discussion that the idiosyncratic component of a firm's total volatility will add noise to the delta-hedged gains. As this noise is firm-specific, the appropriate econometric specification is a *random effects* panel regression.

Consider the regression framework:

$$\text{GAINS}_i(t) \quad = \quad \Omega_0 + \Omega_1 \text{VOL}_m(t) + \epsilon_{i,t} \qquad (19)$$

$$\epsilon_{i,t} = e_{i,t} + u_i \quad (20)$$

The noise term u_i now captures the contribution of the idiosyncratic component of the volatility for firm i.

The regression model is estimated by feasible generalized least squares, and results are presented in Exhibit 4. The estimate of Ω_1 is -0.0263, statistically significant with a p-value of 0.01 (z-statistic of –2.55). These results confirm that individual equity options price a negative market volatility risk premium.

A stronger implication of a market volatility risk premium is that the idiosyncratic component of the firm's volatility, if not priced, must not be related to delta-hedged gains. To examine this hypothesis, we extend the specification of equations (19)-(20) by including the firm's total volatility as an explanatory variable:

$$\text{GAINS}_i(t) = \Omega_0 + \Omega_1 \text{VOL}_m(t) + \Omega_2 \text{VOL}_i(t) + \epsilon_{i,t} \quad (21)$$

$$\epsilon_{i,t} = e_{i,t} + u_i \quad (22)$$

where $VOL_i(t)$ is the i-th firm's return volatility measured over the previous 30 calendar days. If both market and idiosyncratic volatility are priced, then in the augmented regression, Ω_1 should be insignificant and Ω_2 significant. If only market volatility is priced, we should expect contrary results.

Estimating the regression, we find that Ω_1 = –0.0233 with a p-value of 0.03 (z-statistic of –2.15), and Ω_2 = –0.0042 with a p-value of 0.39 (z-statistic of –0.86). Market volatility is significantly correlated with the firm's delta-hedged gains but not the firm's total volatility. The firm's total volatility is not significant even when we drop market volatility from the regression. Thus, the results are consistent with a volatility risk premium where only market volatility risk is priced and idiosyncratic return volatility is unpriced.

What is the economic implication of the estimated coefficient Ω_1? Consider, as an illustration, the call option on AIG. On April 21, 1993, the S&P 500 historical return was 10%. Given the estimate of Ω_1 = –0.0263, the impact of the market volatility risk premium is –0.263% of the stock price level. With AIG stock closing at 125.625, the impact of the market volatility risk premium on the AIG near-money call is \$0.33. The 125 strike call was priced at \$3.82, however, implying a volatility risk premium of 8.65% as a fraction of the call price.

Now consider SPX options. The estimate of Ω_1 is

EXHIBIT 4

Cross-Sectional and Time Series (Panel) Regression—1991–1995

OBS	Ω_1	Ω_2
610	-0.0263 [-2.55]	-
610	-	-0.008 [-1.64]
610	-0.0233 [-2.15]	-0.0042 [-0.86]

z-statistics in brackets.

0.03, and the index level on April 21, 1993, was 444.69. The impact of the volatility risk premium is \$1.33. With the 440 strike call priced at \$8.25, the estimated risk premium is about 16% of the call price. Clearly, the impact of the market volatility risk premium is much higher for index options than for the average stock.

The results overall indicate that market volatility is negatively priced in individual equity options. Our additional insight is that the market volatility risk premium has less of an effect, and idiosyncratic return volatility is unpriced. These results have an important economic implication. If, as has been argued, a negative market volatility risk premium results in greater implied volatilities than realized volatilities, these results explain why this bias is greater for index options than individual equity options. Buyers of individual equity options lose less money, because the market volatility risk premium has less of an impact.

VI. DISCUSSION AND CONCLUSIONS

We have asked whether and to what extent market volatility risk is priced in individual equity options. If market volatility risk is priced, the gains on a delta-hedged option portfolio (long call and short stock) should be negatively correlated with market volatility. A negative volatility risk premium makes an option cost more than it would absent the volatility risk premium. The negative volatility risk premium arises because options serve as a hedge; it is consistent with an inverse correlation between volatility and stock returns.

Our work provides several insights. First, implied volatilities of individual equity options are higher than realized volatilities, although the extent of this bias is substantially smaller than that of the market index.

Second, an investigation of the behavior of delta-hedged gains shows that average delta-hedged gains of individual stocks tend to be negative and are negatively correlated with market volatility. Our results support the hypothesis that the risk premium for market volatility is negative in individual equity markets.

Third, the size of the volatility risk premium indicates a much smaller role for the risk premium for individual stocks than for the market index. Market volatility risk has less of an impact on the losses on delta-hedged gains, and idiosyncratic volatility risk is not priced. Our empirical results provide a perspective on why the bias between implied and realized volatilities is smaller for individual equity options. Overall, our findings are consistent with investor risk aversion with respect to a market component of risk.

Why is the impact of the market volatility risk premium smaller on average on individual equity options than on the index option? First, an idiosyncratic volatility component may mean that changes in market volatility result in a smaller impact on the individual firm's total volatility. In our data set, we find some evidence to support this hypothesis. For instance, when changes in individual volatility are regressed on changes in market volatility, we find the sensitivity coefficient is generally smaller than unity. For this reason, the volatility risk premium has a commensurately smaller effect on individual equity options.

Second, the index option may be more sensitive to market volatility, because market volatility may also affect the pricing of other risks. (Pan [2002], for example, shows that market volatility affects the intensity of jumps and thus the pricing of jump risk.) Given the low negative risk-neutral skewness in individual equities, jump risk appears to be a less important consideration for individual equities.

Although we focus exclusively on the effect of a market volatility risk premium in the time series, much remains to be learned of its effect across a cross-section of individual stock options. Our results regarding a market volatility risk premium suggest that some of the cross-sectional variation in the pricing of individual equity options can be explained by sensitivity of the individual firm's volatility to market volatility. The question of whether the market price of volatility risk is the same across stocks requires estimating individual volatility sensitivities and is quite involved.

If research on cross-sectional differences across stocks is any indication, an investigation of cross-sectional differences among stock options will likely enhance our understanding of how risk-averse investors price derivative assets.

ENDNOTE

The authors thank their colleagues, Stephen Figlewski, Steve Heston, Nengjiu Ju, Dilip Madan, and N. Prabhala, for several helpful suggestions.

REFERENCES

Bakshi, G., C. Cao, and Z. Chen. "Empirical Performance of Alternative Option Pricing Models." *Journal of Finance*, 52 (1997), pp. 2003-2049.

Bakshi, G., and N. Kapadia. "Delta-Hedged Gains and the Negative Volatility Risk Premium." *Review of Financial Studies,* 16 (2) (2003), pp. 527-566.

Bakshi, G., N. Kapadia, and D. Madan. "Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options." *Review of Financial Studies*, 16 (1) (2003), pp. 101-143.

Bates, D. "Post-87 Crash Fears in S&P 500 Futures Options." *Journal of Econometrics,* 94 (2000), pp. 181-238.

Buraschi, A., and J. Jackwerth. "The Price of a Smile: Hedging and Spanning in Option Markets." *Review of Financial Studies,* 14 (2) (2001), pp. 495-527.

French, K., W. Schwert, and R. Stambaugh. "Expected Stock Returns and Volatility." *Journal of Financial Economics,* 19 (1987), pp. 3-29.

Heston, S. "A Closed Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options." *Review of Financial Studies*, 6 (2) (1993), pp. 327-343.

Jackwerth, J., and M. Rubinstein. "Recovering Probability Distributions from Option Prices." *Journal of Finance,* 51 (1996), pp. 1611-1631.

Merton, R. "Option Pricing When Underlying Stock Returns are Discontinuous." *Journal of Financial Economics,* 3 (1976), pp. 125-144.

Pan, J. "The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study." *Journal of Financial Economics*, 63 (2002), pp. 3-50.