

Issues in Hedging Options Positions

**SAIKAT NANDI AND
DANIEL F. WAGGONER**

Nandi is a senior economist and Waggoner is an economist in the financial section of the Atlanta Fed's research department. They thank Lucy Ackert, Jerry Dwyer, and Ed Maberly for helpful comments.

MANY FINANCIAL INSTITUTIONS HOLD NONTRIVIAL AMOUNTS OF DERIVATIVE SECURITIES IN THEIR PORTFOLIOS, AND FREQUENTLY THESE SECURITIES NEED TO BE HEDGED FOR EXTENDED PERIODS OF TIME. OFTEN THE RISK FROM A CHANGE IN VALUE OF A DERIVATIVE SECURITY, ONE WHOSE VALUE DEPENDS ON THE VALUE OF AN UNDERLYING

asset—for example, an option—is hedged by transacting in the underlying securities of the option. Failure to hedge properly can expose an institution to sudden swings in the values of derivatives resulting from large unanticipated changes in the levels or volatilities of the underlying assets. Understanding the basic techniques employed for hedging derivative securities and the advantages and pitfalls of these techniques is therefore of crucial importance to many, including regulators who supervise the financial institutions.

For options, the popular valuation models developed by Black and Scholes (1973) and Merton (1973) indicate that if a certain portfolio is formed consisting of a risky asset, such as a stock, and a call option on that asset (see the glossary for a definition of terms), then the return of the resulting portfolio will be approximately equal to the return on a risk-free asset, at least over short periods of time.¹ This type of portfolio is often called a hedge/replicating portfolio. By properly rebalancing the positions in the underlying asset and the

option, the return on the hedge portfolio can be made to approximate the return of the risk-free asset over longer periods of time. This approach is often referred to as dynamic hedging. However, forming a hedge portfolio and then rebalancing it through time is often problematic in the options market. There are two potential sources of errors: The first is that the option valuation model may not be an adequate characterization of the option prices observed in the market. For example, the Black-Scholes-Merton model says that the implied volatility should not depend on the strike price or the maturity of the option.² In most options markets, though, the implied volatility of an option does depend on the strike price and time to maturity of the option, a phenomenon that runs contrary to the very framework of the Black-Scholes-Merton model itself. The second potential source of error is that many option valuation models, such as the Black-Scholes-Merton model, are developed under the assumption that investors can trade and hedge continuously through time. However, in practice,

investors can rebalance their portfolios only at discrete intervals of time, and investors incur transaction costs at every rebalancing interval in the form of commissions or bid-ask spreads. Rebalancing too frequently can result in prohibitive transaction costs. On the other hand, choosing not to rebalance may mean that the hedge portfolio is no longer close to being optimal, even if the underlying option valuation model is otherwise adequate.

This article examines some strategies often used to offset limitations in the Black-Scholes-Merton model, describing how the risk of existing positions in options can be hedged by trading in the underlying asset or other options. It shows how certain basic hedge parameters such as “deltas,” which are defined and discussed later, are derived given an option pricing model. Subsequently, the discussion notes some of the practical problems that often arise in using the dynamic hedging principles of the Black-Scholes-Merton model and considers how investors and traders try to circumvent some of these problems. Finally, the hedging implications of the simple Black-Scholes-Merton model are tested against certain ad hoc pricing rules that are often used by traders and investors to get around some of the deficiencies of the Black-Scholes-Merton model. The Standard and Poor’s (S&P) 500 index options market, one of the most liquid equity options markets, is used to compare the hedging efficacies of various models. This study suggests that ad hoc rules do not always result in better hedges than a very simple and internally consistent implementation of the Black-Scholes-Merton model.

How Are Option Payoffs Replicated and Deltas Derived?

To hedge an option, or any risky security, one needs to construct a replicating portfolio of other securities, one in which the payoffs of the portfolio exactly match the payoffs of the option. Replicating portfolios can also be used to price options, but this discussion will be limited to

their hedging properties. Before considering the hedging aspects of the Black-Scholes-Merton model, a few simple examples will illustrate how such portfolios are constructed.

One-Period Model.³ The first example is a European call option on a stock, assuming that the stock is currently valued at \$100.⁴ In this example, the option expires in one year and the strike or exercise price is \$100, and the annual risk-free interest rate is 5 percent so that borrowing \$1 today will mean having to pay back \$1.05 one year from now. For simplicity, the assumption is that there are only two possible outcomes when the option expires—the stock price can be either \$120 (an up state) or \$80 (a down state). Note that the value of the call option will be \$20 if the up state occurs and \$0 if the down state occurs as shown below (see Chart 1).⁵

Since there are only two possible states in the future, it is possible to replicate the value of the option in each of these states by forming a portfolio of the stock and a risk-free asset. If Δ shares of the stock are purchased and M dollars are borrowed at the risk-free rate, the stock portion of the portfolio is worth $120 \times \Delta$ in the up state and $80 \times \Delta$ in the down state while $1.05 \times M$ will have to be paid back in either of the states. Thus, to match the value of the portfolio to the value of the option in the two states, it must be the case that

$$120 \times \Delta - 1.05 \times M = 20 \text{ (up state)} \quad (1)$$

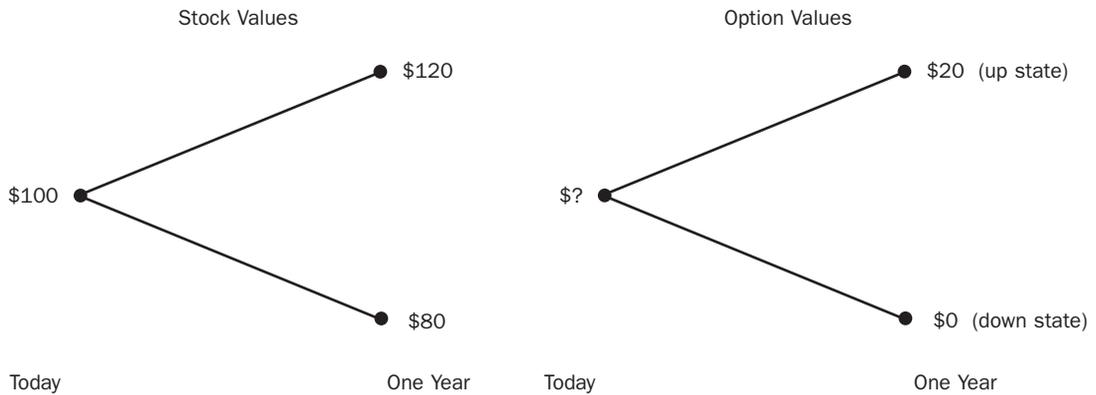
and

$$80 \times \Delta - 1.05 \times M = 0 \text{ (down state)}. \quad (2)$$

Because of the simplicity and tractability of the Black-Scholes-Merton model for valuing options, the model is widely used by options traders and investors.

1. For the purposes of this article, the risk-free asset is a money market account that has no risk of default.
2. Implied volatility in the Black-Scholes-Merton model is the level of volatility that equates the model value of the option to the market price of the option.
3. The fact that results reported in this article have been rounded off from actual values may account for small differences when the computations are recreated.
4. The general principle of hedging discussed here applies not only to stock options but also to interest rate options and currency options. Although not discussed here, deltas for American options can be similarly derived for the example shown here. See Cox and Rubinstein (1985) for American options.
5. Note that the risk-free interest rate of 5 percent lies between the return of 20 percent in the up state and -20 percent in the down state. For example, if the interest rate were above 20 percent, then one would never hold the risky asset because its returns are always dominated by the return on the risk-free asset.

CHART 1 Stock and Option Values in the One-Period Model



The resulting system of two equations with two unknowns (Δ and M) can be easily solved to get $\Delta = 0.5$, and M is approximately 38.10. Therefore, one would need to buy 0.5 shares of the stock and borrow \$38.10 at the risk-free rate in order for the value of the portfolio to be \$20 and \$0 in the up state and down state, respectively. Equivalently, selling 0.5 shares of the stock and lending \$38.10 at the risk-free rate would mean payoffs from that portfolio of $-\$20$ and $\$0$ in the up and down state, respectively, which would completely offset the payoffs from the option in those states.⁶ It is also worth noting that the current value of the option must equal the current value of the portfolio, which is $100 \times \Delta - M = \11.90 .⁷ In other words, a call option on the stock is equivalent to a long position in the stock financed by borrowing at the risk-free rate.

The variable Δ is called the delta of the option. In the previous example, if C_u and C_d denote the values of the call option and S_u and S_d denote the price of the stock in the up and down states, respectively, then it can be verified that $\Delta = (C_u - C_d)/(S_u - S_d)$. The delta of an option reveals how the value of the option is going to change with a change in the stock price. For example, knowing Δ , C_u , and the difference between the stock prices in the up and down state makes it possible to know how much the option is going to be worth in the up state—that is, C_u is also known.

Two-Period Model. A model in which a year from now there are only two possible states of the world is certainly not realistic, but construction of a multiperiod model can alleviate this problem. As for the one-period model, the example for a two-period model assumes a replicating portfolio for a call option on a stock currently valued at \$100 with a

strike price of \$100 and which expires in a year. However, the year is divided into two six-month periods and the value of the stock can either increase or decrease by 10 percent in each period. The semiannual risk-free interest rate is 2.47 percent, which is equivalent to an annual compounded rate of 5 percent. The states of the world for the stock values are given in Chart 2. Given this structure, how does one form a portfolio of the stock and the risk-free asset to replicate the option? The calculation is similar to the one above except that it is done recursively, starting one period before the option expires and working backward to find the current position.

In the case in which the value of the stock over the first six months increases by 10 percent to \$110 (that is, the up state six months from now), the value of the option in the up state is found by forming a replicating portfolio containing Δ_u shares of the stock financed by borrowing M_u dollars at the risk-free rate. Over the next six months, the value of the stock can either increase another 10 percent to \$121 or decline 10 percent to \$99, so that the option at expiration will be worth either \$21 or \$0. Since the replicating portfolio has to match the values of the option, regardless of whether the stock price is \$121 or \$99, the following two equations must be satisfied:

$$121 \times \Delta_u - 1.0247 \times M_u = 21 \quad (3)$$

and

$$99 \times \Delta_u - 1.0247 \times M_u = 0. \quad (4)$$

Solving these equations results in $\Delta_u = 0.9545$ and $M_u = 92.22$. Thus the value of the replicating portfolio is $110 \times \Delta_u - M_u = \12.78 . If, instead, six months

CHART 2
Stock Values in the Two-Period Model

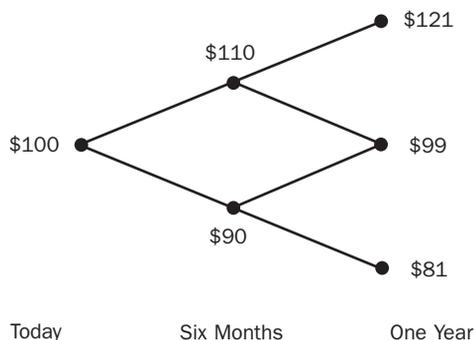
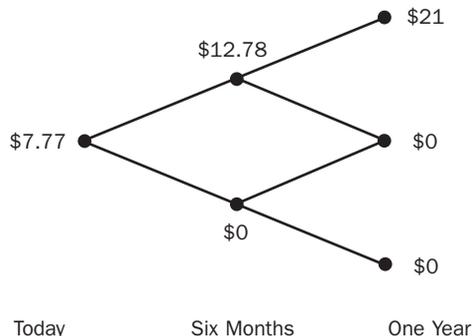


CHART 3
Option Values in the Two-Period Model



from now the stock declines 10 percent in value, to \$90 (the down state), the stock price at the expiration of the option will either be \$99 or \$81, which is always less than the exercise price. Thus the option is worthless a year from now if the down state is realized six months from now, and consequently the value of the option in the down state is zero. Given the two possible values of the option six months from now, it is now possible to derive the number of shares of the stock that one needs to buy and the amount necessary to borrow to replicate the option payoffs in the up and down states six months from now. Since the option is worth \$12.78 and \$0 in the up and down states, respectively, it follows that

$$110 \times \Delta - 1.0247 \times M = 12.78, \quad (5)$$

and

$$90 \times \Delta - 1.0247 \times M = 0. \quad (6)$$

Solving the above equations results in $\Delta = 0.6389$ and $M = 56.11$. Thus the value of the option today is $100 \times \Delta - M = \7.77 . The values of the option are shown graphically in Chart 3.

A feature of this replicating portfolio is that it is always self-financing; once it is set up, no further external cash inflows or outflows are required in the future. For example, if the replicating portfolio is set up by borrowing \$56.11 and buying 0.6389 shares of the stock and in six months the up state is realized,

the initial portfolio is liquidated. The sale of the 0.6389 shares of stock at \$110 per share nets \$70.28. Repaying the loan with interest, which amounts to \$57.50, leaves \$12.78. The new replicating portfolio requires borrowing \$92.22. Combining this amount with the proceeds of \$12.78 gives \$105, which is exactly enough to buy the required 0.9545 (Δ_u) shares of stock at \$110 per share. Replicating portfolios always have this property: liquidating the current portfolio nets exactly enough money to form the next portfolio. Thus the portfolio can be set up today, rebalanced at the end of each period with no infusions of external cash, and at expiration should match the payoff of the option, no matter which states of the world occur.

In the replicating portfolio presented above, the option expires either one or two periods from now, but the same principle applies for any number of periods. Given that there are only two possible states over each period, a self-financing replicating portfolio can be formed at each date and state by trading in the stock and a risk-free asset. As the number of periods increases, the individual periods get shorter so that more and more possible states of the world exist at expiration. In the limit, continuums of possible states and periods exist so that the portfolio will have to be continuously rebalanced. The Black-Scholes-Merton model is the limiting case of these models with a limited number of periods.

6. In other words, a long position in one unit of the option can be hedged by holding a short position in 0.5 shares of the stock and lending \$38.10 at the risk-free rate: the value of the total position is \$0 in both states.
7. If the current value of the option were higher/lower than the value of the replicating portfolio, then an investment strategy could be designed by selling/buying the option and forming the replicating portfolio such that one will always make money at no risk, often called an arbitrage opportunity.

Thus the Black-Scholes-Merton model must assume that investors can trade, or rebalance, continuously through time.⁸ Another assumption of the Black-Scholes-Merton model concerns the volatility of the stock returns over each time period. Volatility is related to the up and down movements in the limited-period models. The Black-Scholes-Merton model assumes that the volatility of the stock returns is either constant or varies in such a way that future volatilities can be anticipated on the basis of current information.⁹

Although the continuous trading assumption may seem unrealistic, the Black-Scholes-Merton model nevertheless provides traders and investors with a very convenient formula in which all the input variables but one are observable. The only unobservable input variable is the implied volatility, that is, the average expected volatility of the asset returns until the option expires. A reasonable guess about the expected future volatility is not very difficult, however, because one can estimate the prevalent volatility from the history of asset prices to the present time. From a trader's or investor's perspective, using the Black-Scholes-Merton formula, then, requires only guessing the implied volatility.¹⁰ A more sophisticated option pricing model, in contrast, may require the trader to guess values of model variables more difficult to obtain in real time, such as the speed of mean reversion of volatility and others. In fact, the simplicity of the Black-Scholes-Merton model largely explains its widespread use regardless of some of its glaring biases from a theoretical perspective. Despite the Black-Scholes-Merton model's very convenient pricing formula, it seems to have serious constraints: it does not allow forming a self-financing replicating portfolio with the provision that one can trade only at discrete intervals of time with nonnegligible transaction costs such as commissions or bid-ask spreads.

Delta Hedging under the Black-Scholes-Merton Model. Considering a European call option on a nondividend paying stock will illustrate some of the shortcomings of the Black-Scholes-Merton model.¹¹ This example assumes that the option has a strike price of \$100 and expires in 100 days; that the current stock price is \$100 and the implied volatility is 15 percent annually; and that the current annual risk-free rate, continuously compounded, is 5 percent. If 100 call options have been written (100 options typically constitute an options contract), a delta-neutral portfolio will have to be formed to hedge exposure to stock price movements. A delta-neutral portfolio is one that is insensitive to small changes in the price of the underlying stock. Using the Black-Scholes-Merton option valuation formula

given in Box 1, the value of each option is approximately \$3.8375, so that \$383.75 is received by selling or writing the option. Since the portfolio should be self-financing, the proceeds from the options are invested in the stock and risk-free asset. Thus \$383.75 is invested in a portfolio of N shares of the stock and in M dollars of the risk-free asset.

Let Δ denote the delta of the option and, in accordance with the formula for Δ for the Black-Scholes-Merton model given in Box 1, $\Delta = 0.5846$. The delta of the total position (option, stock, and risk-free asset) is a linear combination of the deltas of the options, the stock, and the risk-free asset. The delta of a long (short) position in the option is Δ ($-\Delta$), the delta of a long (short) position in the stock is 1 (-1), and the delta of the risk-free asset is zero. As 100 options have been sold and N shares have been bought, the delta of the portfolio is $-100 \times \Delta + N$.

In order for the portfolio to be delta-neutral, the following equation must be satisfied:

$$-100 \times \Delta + N = 0. \quad (7)$$

Similarly, for the portfolio to be self-financing, it has to be the case that

$$N \times 100 + M = 383.75. \quad (8)$$

In solving the two equations above for N and M , $N = \Delta \times 100 = 58.46$ and $M = -5,462.25$. Thus 100 options have been sold for a total of \$383.75, 58.46 units of the share have been bought, and \$5,462.25 has been borrowed at an annual interest rate of 5 percent. The total value of the portfolio is zero when it is formed because the portfolio is self-financing. What happens, though, to the portfolio value on the next trading day for three different levels of the stock prices? Borrowing \$5,462.25 has incurred interest charges of approximately $\$5,462.25 \times 0.05/365.0 = \0.748 . Thus the value of the portfolio on the next day (denoted as $t + 1$) is

$$V(t + 1) = 58.46 \times S(t + 1) - 100 \times C(t + 1) - (5,462.25 + 0.748), \quad (9)$$

where $S(t + 1)$ and $C(t + 1)$ denote the values of the stock and the call option, respectively, on the next day. Table 1 gives the value of the option and thereby the value of the delta-neutral portfolio for various values of the stock price, assuming that everything else (including the implied volatility) is the same.

The value of the delta-neutral portfolio is not zero in any of these cases, even though in one the stock price did not change from its initial value of \$100. The reason is that the delta has been derived from a

model that assumes continuous trading and thus requires continuous rebalancing for the delta-neutral portfolio to retain its original value. Transactions costs, like broker commissions and margin requirements, would further deteriorate the performance of the delta-neutral portfolio.¹²

Other Dynamic Hedging Procedures Using the Black-Scholes-Merton Model. The previous example assumed that the underlying Black-Scholes-Merton model generated the option prices so that the implied volatility was the same on both days. However, in reality the implied volatility is not constant but changes through time in almost all options markets. The following example demonstrates the outcome if the implied volatility changes on the next day, assuming that the implied volatility on the next day ($t + 1$) is 15.5 percent, 15 percent, and 14.5 percent, corresponding to three different stock prices of \$99, \$100, and \$101. The fluctuation of implied volatility suggested here corresponds to stock price, increasing as the stock price goes down and decreasing as it goes up—a feature of many equity and stock index options markets. Table 2 shows the values of the portfolio corresponding to three different levels of stock prices and implied volatilities.

Thus, with a change in the implied volatility of around 0.5 percent (frequently observed in options markets), the hedging performance of the Black-Scholes-Merton model deteriorates quite sharply. The hedge portfolios constructed on the previous day are quite poor primarily because the model's assumption of constant variance is violated. Extensive academic literature documents how implied volatilities in the options market change through time (Rubinstein 1994; Bates 1996; and many others).¹³ Further, volatility often varies in ways that cannot always be predicted with current information. How could traders or investors set up hedge portfolios that would account for the random variation in volatilities? One alternative is to derive

TABLE 1
The Delta-Neutral Portfolio on the Next Day with No Change in Implied Volatility

Stock Price	Option Price	Portfolio Value
\$ 99	\$3.26	-\$0.96
\$100	\$3.82	\$1.53
\$101	\$4.42	-\$0.84

TABLE 2
The Delta-Neutral Portfolio on the Next Day When Implied Volatility Changes

Stock Price	Implied Volatility (Percent)	Portfolio Value
\$ 99	15.5	-\$11.26
\$100	15.0	\$ 1.50
\$101	14.5	\$ 9.06

the hedge portfolio from a more sophisticated (and more complex) option pricing model such as a stochastic volatility model (to be discussed later). However, estimating and implementing such a model can be difficult for an average trader or investor. Practitioners may be better served by finding ways to circumvent the hedging deficiencies of the Black-Scholes-Merton model stemming from implied volatilities that change through time but sticking to the model as much as possible.

One way to get around the problem of time-varying volatility that occurs with the Black-Scholes-Merton model is to form a hedge portfolio that is insensitive to both the changes in the price of the underlying asset and its volatility. The sensitivity of an option price with respect to the volatility is often referred to as vega. In order to hedge against changes in both the asset price and volatility, one can form a portfolio that is delta-neutral as well as

8. This replication with continuous trading is possible due to a special property known as the martingale representation property of Brownian motions (see Harrison and Pliska 1981).
9. However, with continuous trading, one can form a self-financing portfolio by trading in the stock and the risk-free asset even if the volatility of the stock is random. All that is needed is that the Brownian motions driving the stock price and the volatility are perfectly correlated (see Heston and Nandi forthcoming).
10. Given the existence of multiple implied volatilities from different options (on the same asset), this task is a little more complicated.
11. If the stock pays dividends, then the present value of the dividends that are to be paid during the life of the option must be subtracted from the current asset price; the resulting asset price is used in the option pricing formula.
12. It is also worth noting that the portfolio is not self-financing on the next day because rebalancing would incur an external cash flow in each of the three states.
13. One can also go to the Web site www.cboe.com/tools/historical/vix1986.txt to see the daily history of the implied volatility index on the Standard and Poor's 100, called the VIX. VIX captures the implied volatilities of certain near-the-money options on the Standard and Poor's 100 index (ticker symbol, OEX).

Black-Scholes Price and Deltas

The Black-Scholes-Merton formula gives the current value of a European call/put option in terms of (a) $S(t)$, the price of the underlying asset; (b) K , the strike or exercise price; (c) τ , the time to maturity of the option; (d) $r(\tau)$, the risk-free rate or the equivalent yield of a zero-coupon bond (that matures at the same time as the option); and (e) σ , the square root of the average per period (for example, daily) variance of the returns of the underlying asset that will prevail until the option expires.¹ Assuming that the underlying asset does not pay any dividends until the option expires, the call and put values are at time t .

$$C(t) = S(t) N(d1) - K \exp[-r(\tau)\tau] N(d2), \quad (B1)$$

and

$$P(t) = K \exp[-r(\tau)\tau] N(-d2) - S(t) N(-d1), \quad (B2)$$

where $N()$ is the standard normal distribution function and

$$d1 = \{\ln(S/K) + [r(\tau) + 0.5\sigma^2]\tau\} / \sigma\sqrt{\tau} \quad (B3)$$

and

$$d2 = d1 - \sigma\sqrt{\tau}. \quad (B4)$$

(The tables for computing the function are found in almost all basic statistics books.) If the underlying asset pays known dividends at discrete dates until the option expires, then the present value of the dividends must be subtracted from the asset price to substitute for $S(t)$ in the above formulas.² Of the above-mentioned variables that are required as inputs to the Black-Scholes-Merton formula, only σ is not readily observable.

The delta of the option is the partial derivative of the option price with respect to the asset price, that is, dC/dS for call options and dP/dS for put options. An important property of the Black-Scholes-Merton formula is that the option price is homogeneous of degree 1 in the asset price and the strike price. Hence it follows from Euler's theorem on homogeneous functions (see Varian 1984) that the delta of the call option is $N(d1)$ and that of the put option is $N(d1) - 1$.

The vega of a call or put option is $dC/d\sigma$ or $dP/d\sigma$. Hull (1997, 329) gives the formula for vega in terms of the same variables that appear in the valuation formula.

1. Actually the Black-Scholes (1973) model assumes that the risk-free rate is constant. However, Merton (1973) shows that even if interest rates are random, the appropriate interest rate to use in the Black-Scholes formula for a stock option is the yield of a zero-coupon bond that expires at the same time as the option. In that case, the simple Black-Scholes (1973) formula serves as an extremely good approximation because the volatility of interest rates is relatively low compared with the volatility of the underlying stock.
2. The corresponding exact valuation formula for American put options (or call options on dividend paying assets) and deltas are not known explicitly. However, there are good analytical approximations as in Carr (1998), Ju (1998), and, Huang, Subrahmanyam, and Yu (1996).

vega-neutral. The formation of such a portfolio is indeed ad hoc: in fact, it is theoretically inconsistent because under the Black-Scholes-Merton model volatility is constant (or deterministic) and therefore does not need to be hedged. Forming a delta-vega-neutral portfolio would require trading two options, the underlying asset and the risk-free asset.

Adding to the previous example, in which an option contract has been sold (with 100 days to expire) and in which all other variables such as the stock price and the strike price are the same as before, $N2$ units of a second option, $N3$ units of the stock, and M dollars of the risk-free asset are

required. The current values of the first and second option are denoted as $C(1)$ and $C(2)$, respectively, whereas the current stock price is denoted as $S(t)$. Since the second option can be chosen freely, an option of the same strike (\$100) but a maturity of 150 days is selected. Given these, $C(1) = \$3.8375$ and $C(2) = \$4.898$. The current deltas of the two options are denoted as $\Delta(1)$ and $\Delta(2)$, and the vegas, as $\text{vega}(1)$ and $\text{vega}(2)$ (see Hull 1997 for the formula for vega).

For the delta of the portfolio to be zero, it is necessary that

$$-100 \times \Delta(1) + N2 \times \Delta(2) + N3 = 0. \quad (10)$$

For the vega of the portfolio to be zero, it is necessary that

$$-100 \times \text{vega}(1) + N2 \times \text{vega}(2) = 0.14 \quad (11)$$

For the portfolio to be self-financing, it is necessary that

$$\begin{aligned} -100 \times C(1) + N2 \times C(2) \\ + N3 \times S(t) - M = 0. \end{aligned} \quad (12)$$

Solving the equations in this system of three equations with three unknowns ($N2$, $N3$, and M) shows that $N2 = 82.59$, $N3 = 8.64$, and $M = \$884.96$. Thus 82.59 units of the second option and 8.64 units of the stock must be bought, and \$884.96 must be borrowed at the risk-free rate. Table 3 shows the value of the delta-vega-neutral portfolio on the next day. The terms $C(1)_{t+1}$ and $C(2)_{t+1}$ denote the prices of the first and second option on the next day.

The delta-vega-neutral hedge portfolio performs much better than a delta-neutral hedge portfolio that uses just one option, especially if the implied volatilities change. The only disadvantage in using this kind of hedge is that the portfolio requires two options, and options markets tend to be less liquid than the market on an underlying asset, such as a stock. On average, options have much higher bid-ask spreads (relative to their transaction prices) than those on an underlying asset such as a stock. Using a second option to hedge the volatility risk therefore could increase transaction costs, especially for a retail investor.

Similar to delta-vega hedging is what is known as delta-gamma hedging. The gamma of an option measures the rate of change of its delta with respect to a change in the price of the underlying asset. The more the delta of the option changes with the asset price, the more a portfolio will have to be rebalanced to remain delta-neutral. The purpose of delta-gamma hedging is to create a portfolio that is both delta-neutral and gamma-neutral. Thus, ceteris paribus, the amount of rebalancing required in a delta-gamma-neutral portfolio would tend to be lower than that in a delta-neutral portfolio over short periods of time, and lower rebalancing could be used to offset higher transactions costs. Constructing a delta-gamma-neutral portfolio also requires two options; the number of units of the second option can be found by solving a similar set of equations to those applied to the delta-vega-neutral portfolio discussed previously. A delta-vega-gamma-neutral portfolio can also be created,

TABLE 3
The Delta-Vega-Neutral Portfolio on the Next Day When Implied Volatility Changes

Stock Price	Implied Volatility (Percent)	$C(1)_{t+1}$	$C(2)_{t+1}$	Portfolio Value
\$ 99	15.5	\$ 3.36	\$ 4.42	\$ -0.30
\$100	15	\$ 3.81	\$ 4.88	\$ 0.51
\$101	14.5	\$ 4.32	\$ 5.38	\$ -0.34

but forming such a portfolio requires positions in three options.

The hedging problems discussed thus far fall under the rubric of dynamic hedging in that they require a portfolio formed of the underlying asset and a risk-free asset or options that must be rebalanced through time. Since the number of units of the underlying asset and the risk-free asset or other options are derived from an option pricing model, such as the Black-Scholes-Merton model, the formation of the hedge portfolio is prone to model misspecifications; that is, the underlying options valuation model is not consistent with the option prices observed in the market. An alternative to dynamic hedging is static hedging in which a portfolio is formed as of today and requires no further trading in the underlying asset and options.

Let S , K , P , and C denote the underlying asset price, strike price, put price, and call price, respectively. (Note that both the put and call have the same strike price.) If r and τ denote the risk-free rate and time to expiration, then the put-call parity relationship for European options states that the following has to hold exactly at any given point in time (in the absence of transactions costs):

$$P = C - S + Ke^{-r\tau}. \quad (13)$$

Thus, to replicate the payoff of a put option with the strike price, K , and time to maturity, τ , a synthetic portfolio must be constructed containing a call option of the same strike and maturity as that of the put, a short sell of the asset, and a long position on K units of a discount bond (that pays off \$1 at maturity) that matures at the same time as the options. Once the synthetic portfolio has been set up for the put option, rebalancing is no longer necessary because the price of the put option is identical to that of the synthetic portfolio if put-call parity is to be preserved. Since the put-call parity relationship is

14. The vega of a portfolio of options is a linear combination of the vegas of the individual options, and the vega of the underlying asset is zero. Vega(1) = 20.41; vega(2) = 24.71; $\Delta(1) = 0.585$; $\Delta(2) = 0.603$.

independent of any option valuation model, static hedging may seem to be the preferable path. However, static hedging is also prone to some of the same drawbacks that occur when options are hedged with options—namely, that options markets are relatively illiquid, and the second option may not be available in the right quantity. For example, in the Standard and Poor's 500 index options, a market maker may have to satisfy huge buy order flows in deep out-of-the-money put options—those with strike prices substantially below the current S&P 500 level—from institutional investors who want to hedge their positions against sharp downturns in the index. However, the volume of deep-in-the-money call options that would be required in the hedge/replicating portfolio (as per put-call parity) is relatively low, and hedging deep-out-of-the-money puts via deep-in-the-money calls may not be readily feasible.

Static hedging has often been advocated as a useful tool for certain types of exotic options known as barrier options.¹⁵ Barrier options tend to have regions of very high gammas; that is, the delta changes very rapidly and thus requires frequent rebalancing in certain regions (for example, if the asset price is close to the barrier). Dynamic hedging may therefore turn out to be quite difficult and costly for barrier options. Nevertheless, liquidity issues concerning static hedging discussed previously also apply to barrier options. A further difficulty is that some options needed as part of the static hedge portfolios for barrier options may not be traded at all, so close substitutes must be chosen. In hedging exotic options such as barrier options, a trade-off between the pros and cons of static and dynamic hedging is thus inevitable.

Smile, Smirk, and Hedge. Because of its simplicity (traders have to guess only one unobservable variable—the average expected volatility of the underlying asset over the life of the option) the Black-Scholes-Merton model continues to be very popular with most traders. However, from a theoretical perspective, the model always exhibits certain biases. One very prevalent and widely documented bias is that the implied volatilities in the Black-Scholes-Merton model depend on the strike price and maturity of an option. Chart 4 shows the implied volatilities in the Standard and Poor's 500 index options for call options of different strike prices on December 21, 1995, with twenty-eight and fifty-six days to maturity. The implied volatilities in the Standard and Poor's 500 index options market tend to decrease as the strike price increases; this pattern is sometimes referred to as a volatility smirk. Similarly, in some other options markets, such as the currency options market, the implied

volatilities decrease initially as the strike price increases and then increase a little—a U-shaped pattern often referred to as a smile. Chart 4 also makes it apparent that for options of the same strike price, implied volatility differs depending on the maturity of the option. For example, if the strike price is \$570, the implied volatility of the option with twenty-eight days to maturity is 18.7 percent whereas the implied volatility of the option with fifty-six days to maturity is 16.7 percent. Such variations in implied volatilities across strike prices and maturities are inconsistent with the basic premise of the Black-Scholes-Merton model, which accommodates only one implied volatility irrespective of strike prices and maturities. Before examining the hedging implications of this bias, it is important to understand what could possibly be causing such a phenomenon for index options.

One possibility for the existence of the smirk pattern in implied volatilities is that the options market expects the Standard and Poor's 500 index to go down with a higher probability than that suggested by the statistical distribution postulated for the returns of the index in the Black-Scholes-Merton model. As a result, the market would put a higher price on an out-of-the-money put than would the Black-Scholes-Merton model. Since option prices (both puts and calls) under Black-Scholes-Merton increase as volatility increases, the implied volatility using the Black-Scholes-Merton model would be higher than it would otherwise be. In fact, if the distribution of the returns of the underlying asset is seen as embedded in a cross section of option prices with different strike prices (see Jackwerth and Rubinstein 1996), the distribution appears to be one in which, given today's index level, the probability of negative returns in the future is higher than the probability of positive returns of equal magnitude. Such distributions are said to be skewed to the left.¹⁶ In contrast, the statistical distribution that drives the returns of an underlying asset under the Black-Scholes-Merton model is Gaussian/normal, which does not involve skewness. In other words, given today's index level, the probability of positive returns is the same as the probability of negative returns of equal magnitude.

Is it possible to get such negatively skewed distributions under alternative assumptions of the statistical process that generates returns? It turns out that allowing for future changes in volatility to be random and allowing volatility to be negatively correlated with the returns of the underlying asset can generate negatively skewed distributions of the returns of the underlying asset.¹⁷ Indeed, option pricing models have been developed in which the

volatility of the underlying asset varies randomly through time and is correlated with the returns of the underlying asset. One class of such models, known as implied binomial tree/deterministic volatility models, was first proposed by Dupire (1994), Derman and Kani (1994), and Rubinstein (1994). In these models the current volatility (sometimes known as local volatility) is a function of the current asset price and time, unlike in the Black-Scholes-Merton model, in which volatility is constant through time.¹⁸ These models are also known as path-independent time-varying volatility models in that the current volatility does not depend on the history or path of the asset price. In another class of models, sometimes known as path-dependent time-varying volatility models, the current volatility is the function of the entire history of asset prices and not just the current asset price.¹⁹

Testing the hedging efficacy of an option valuation model often involves measuring the errors incurred in replicating the option with the prescribed replicating portfolio of the model. In other words, the replicating portfolio is formed today, and at a future time the value of the replicating portfolio is compared with the option price observed in the market as of that time. In empirical tests of path-independent time-varying volatility models, Dumas, Fleming, and Whaley (1998) show that in the Standard and Poor's 500 index options market the replication errors of delta-neutral portfolios of path-independent volatility models are greater than those of the very simple Black-Scholes-Merton model. In fact, in terms of replication errors of delta-neutral portfolios, a very simple implementation of the model also appears to dominate an ad hoc variation of the model that uses a separate implied volatility for each option to fit to the smile/smirk curve. The Black-Scholes-Merton model proves more useful for hedging despite the

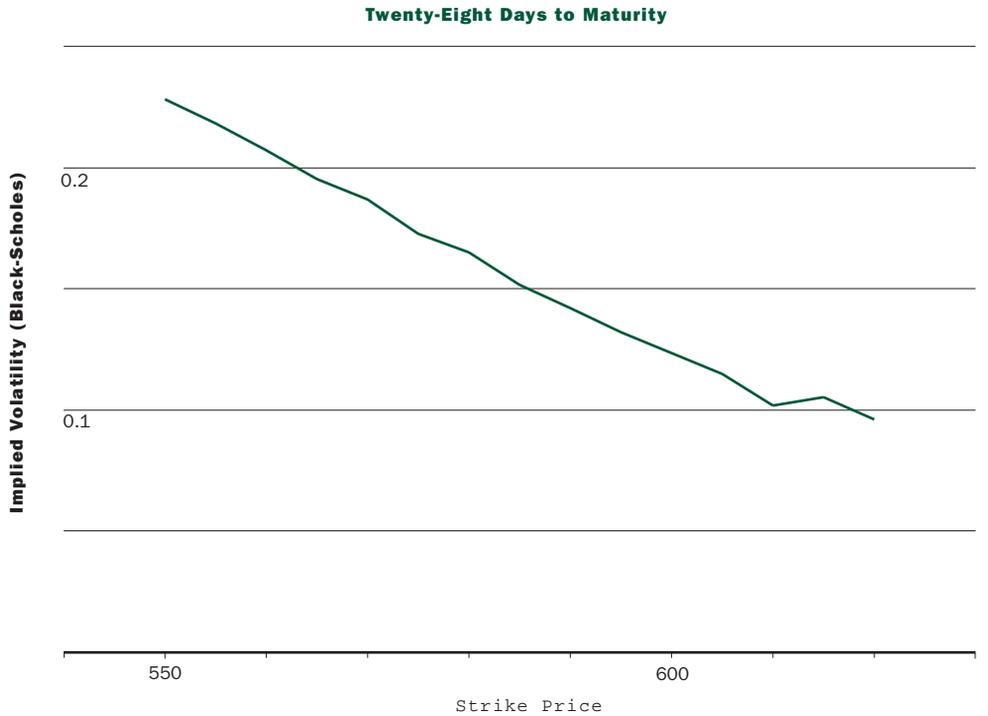
fact that in terms of predicting option prices (that is, computing option prices out-of-sample) it is dominated by the ad hoc rule and the time-varying path-independent volatility model.

Why is it more useful? As discussed above, the hedge ratio, or the delta, which measures the rate of the change in option price with respect to the change in the price of the underlying asset, is an important consideration. If a replicating/hedge portfolio (from an option pricing model) is formed to replicate the value of the option at the next period, it can be shown that to a large extent the hedging/replication error reflects the difference in the pricing or valuation error between the two periods (see Dumas, Fleming, and Whaley 1998). Though one model, model A for example, may result in a lower pricing error (even out-of-sample) than another model, in order for model A to result in lower hedging errors than model B, it could also often be necessary that the change (across two time periods) in valuation error under model A be less than that under model B. More often than not, however, the differences in the valuation errors (across two time periods) between models turn out not to be very significant for most classes of options (that is, options of different strike prices and maturities). In other words, although the Black-Scholes-Merton model exhibits pricing biases, as long as these biases remain relatively stable through time, its hedging performance can be better than the performance of a more complex model that can account for many of the biases, especially if the more complex model does not adequately characterize the way asset prices evolve over time.

Hedging with Ad Hoc Models. How do traders or investors who routinely use the Black-Scholes-Merton model to arrive at hedge ratios/deltas use the model, despite the fact that patterns in implied volatilities across options of different strike prices

15. An example of a barrier option is a down-and-out call option in which a regular call option gets knocked out; that is, it ceases to exist if the asset price hits a certain preset level.
16. The distribution that is skewed is the risk-neutral distribution of asset returns (see Nandi 1998 for risk-neutral probabilities/distributions) and not necessarily the actual distribution of asset returns.
17. Negative correlation implies that lower returns are associated with higher volatility. As a result, the lower or left tail of the distribution spreads out when returns go down, generating negative skewness. This negative correlation is often referred to as the leverage effect (Black 1976; Christie 1982) in equities. One possible explanation for this effect is that as the stock price goes down, the amount of leverage (ratio of debt to equity) goes up, thus making the stock more risky and thereby increasing volatility. An argument against this explanation is that the negative correlation can be observed for stocks of corporations that do not have any debt in their capital structure.
18. Since the future level of the asset price is unknown, the future local volatility is also not known, and, strictly speaking, unlike in the Black-Scholes-Merton model, volatility is not deterministic in these models.
19. See Heston (1993) and Heston and Nandi (forthcoming) for option pricing models with path-dependent volatility models in continuous and discrete time, respectively. These models are sometimes known as continuous time stochastic volatility and discrete-time GARCH models, respectively. Continuous time models are very difficult to implement due to the fact that volatility is unobservable given the history of asset prices.

CHART 4 Implied Volatilities of Call Options



Note: The chart shows the implied volatilities from Standard and Poor's 500 call options of different strike prices on December 21, 1995. The Standard and Poor's 500 index level was at approximately 610.

Parameter Estimation

The Black-Scholes-Merton-2 version of the model uses a procedure called nonlinear least squares (NLS) to estimate a single implied volatility across all options each Wednesday. The NLS procedure minimizes the squared errors between the market option prices and model option prices. The difference between the model price (given an implied volatility, σ) and the observed market price of the option is denoted by $e_i(\sigma)$. As mentioned in Box 1, the midpoint of the bid-ask quote is used for the

observed market price of the option. Thus the criterion function minimized at each t (over σ) is

$$\sum_{i=1}^{N_t} e_i(\sigma)^2,$$

where N_t is the number of sampled bid-ask quotes on day t . In essence, this procedure attempts to find a single implied volatility that minimizes the squared pricing errors of the model.

and maturities are inconsistent with the model? As it turns out, such traders or market makers often use certain theoretically ad hoc variations of the basic Black-Scholes-Merton model to circumvent its biases. Such ad hoc variations allow the implied volatilities input to the Black-Scholes-Merton model to differ across strike prices and maturities. Using a separate implied volatility for each option is inconsistent with the basic theoretical underpinning of the Black-Scholes-Merton model, but it is a common practice among traders and market makers in certain options exchanges (Dumas, Fleming, and Whaley 1998). In the course of implementing such ad hoc variations, options traders or investors can be thought of as using the Black-Scholes-Merton model as a translation device to express their opinion on a more complicated distribution of asset returns than the Gaussian distribution that underlies the Black-Scholes-Merton model.

Ad hoc variations of the basic Black-Scholes-Merton model, depending on the way they are designed, may result in prices that better match observed market prices. But do they necessarily result in better hedging performance? Four versions of the Black-Scholes-Merton model that differ from one another in terms of fitting a cross section of option prices (in-sample errors) and also in predicting option prices (out-of-sample errors) will be presented; these examples illustrate that the differences between the models in terms of hedging/replication errors are not as significant as the differences in valuation errors for most options. In fact, if the models are ranked in terms of the replication errors of the delta-neutral portfolios, the ranking could prove different than when the models are ranked in terms of valuation errors.

There are many different ways in which a trader or investor can input a value for volatility in the Black-Scholes-Merton formula for computing the delta of an option. The Black-Scholes-Merton model assumes that the volatility of an asset's returns is constant through time. However, an investor trying to use the model in the real world is not constrained to hold the volatility constant and can periodically estimate volatility from past observations of asset prices. As an alternative to using the historical data, a single implied volatility for all options (of different strikes and maturities every day) can be estimated that minimizes a criterion function involving the squared price differentials between model prices and the observed prices in the market (see Box 2 for details).

This approach results in a single implied volatility for all options every day. On the other hand, implied volatility can be based on observation of a particular option so that a different implied volatility exists for each option. As an alternative to using the exact implied volatility for each option, a procedure that “merely smoothes Black/Scholes implied volatilities across exercise prices and times to expiration” is used by some options market makers at the Chicago Board Options Exchange (CBOE) (Dumas, Fleming, and Whaley 1998). For example, given that the shape of the smirk in implied volatilities resembles a parabola, one can choose the implied volatility to be a function of the strike price and the square of the strike price. However, implied volatilities differ across maturities even for the same strike price. Thus the time to maturity—and possibly the square of the time to maturity—can also be included in the function. The equation below is used in Dumas, Fleming, and Whaley (1998).

$$\sigma(K, \tau) = a_0 + a_1K + a_2K^2 + a_3\tau + a_4\tau^2 + a_5K\tau, \quad (14)$$

where K is the strike price and τ is the time to maturity of the option. Since the implied volatility $\sigma(K, \tau)$ is observable for each K and τ , one can use the above equation as an ordinary least squares (OLS) regression of the implied volatilities on the various right-hand variables to get the coefficients a_0, a_1, a_2 , and so on. These coefficients provide an estimated implied volatility for each option.²⁰ To summarize, one can use the Black-Scholes-Merton model to arrive at the delta in four different ways: (a) compute the delta with volatility estimated from historical prices, (b) compute the delta using a single implied volatility that is common across all options, (c) compute the delta using the exact implied volatility for each option, and (d) compute the delta using an estimated implied volatility for each option that fits to the shape of the smirk across strike prices and time to maturities.

Of the four different versions of the Black-Scholes-Merton discussed above, the two that allow implied volatilities to differ across options of different strike prices and maturities are indeed ad hoc. The other two versions that result in a single implied volatility across all strikes and maturities are much less ad hoc. Implementing the four different versions of the Black-Scholes-Merton model in the Standard and Poor's 500 index options makes it possible to explore the differences in hedging errors produced by these approaches.

The market for Standard and Poor's 500 index options is the second most active index options market in the United States, and in terms of open interest in options it is the largest. It is also one of the most liquid options markets.²¹ These models test data for the time period from January 5, 1994, to October 19, 1994.²² Box 3 gives a detailed description of the options data used for the empirical tests. The replicating/hedge portfolios are formed on day t from the first bid-ask quote in that option after 2:30 P.M. (central standard time). The portfolio is liquidated on one of the following days— $t + 1$, $t + 3$, or $t + 5$.²³ The hedging error for each version of the Black-Scholes-Merton model is the difference between the value of the replicating portfolio and the option price (measured as the midpoint of the bid-ask prices) at the time of the liquidation.

The first panel of Table 4 shows the mean absolute hedging errors (for the whole sample and across all options) of the four versions of the Black-Scholes-Merton (BSM) model.²⁴ Black-Scholes-Merton-1 is the version of the model in which volatility is computed from the last sixty days of closing Standard

and Poor's 500 index levels. Black-Scholes-Merton-2 is the version of the model in which a single implied volatility is estimated for all options each day. Ad hoc-1 is the ad hoc version of the Black-Scholes-Merton model in which each option has its own implied volatility each day, and ad hoc-2 is the other ad hoc version, in which the implied volatility (on each day) for each option is estimated via the OLS procedure discussed previously.

The first panel clearly shows that judging models on the basis of hedging/replication errors could be somewhat different from judging them on the basis of valuation errors, as discussed previously; valuation errors could include either in-sample errors that show how well the model values fit market prices or out-of-sample/predictive error.²⁵ For example, ad hoc-2 yields substantially lower prediction errors than the Black-Scholes-Merton-2 version (Heston and Nandi forthcoming) but is the least competitive in terms of hedging errors. On the other hand, the magnitude of hedging errors of ad hoc-1, in which the in-sample valuation errors is essentially zero (as each option is priced exactly), is not very different from that of Black-Scholes-Merton-1. In fact, Black-Scholes-Merton-1, which has the highest in-sample valuation errors (as volatility is not

TABLE 4 Mean Absolute Hedging Errors

	BSM-1	BSM-2	Ad Hoc-1	Ad Hoc-2
Whole Sample, All Options				
One-day	\$0.46	\$0.45	\$0.43	\$0.52
Three-day	\$0.66	\$0.65	\$0.62	\$0.78
Five-day	\$0.98	\$0.94	\$0.87	\$1.07
Far-out-of-the-Money Puts under Forty Days to Maturity				
One-day	\$0.22	\$0.16	\$0.10	\$0.19
Three-day	\$0.23	\$0.19	\$0.20	\$0.26
Five-day	\$0.63	\$0.50	\$0.40	\$0.64
Near-the-Money Calls under Forty Days to Maturity				
One-day	\$0.25	\$0.33	\$0.24	\$0.34
Three-day	\$0.49	\$0.52	\$0.44	\$0.60
Five-day	\$0.98	\$1.08	\$0.90	\$0.83
Near-the-Money Puts Forty to Seventy Days to Maturity				
One-day	\$0.52	\$0.56	\$0.53	\$0.62
Three-day	\$0.74	\$0.76	\$0.77	\$0.91
Five-day	\$1.20	\$1.34	\$1.15	\$1.17

Source: Calculated by the Federal Reserve Bank of Atlanta using data from Standard and Poor's 500 index options market

Data Description

The data set used for hedging is a subset of the tick-by-tick data on the Standard and Poor's 500 options that includes both the bid-ask quotes and the transaction prices; the raw data set is obtained directly from the exchange. The market for Standard and Poor's 500 index options is the second most active index options market in the United States, and in terms of open interest in options it is the largest. It is also easier to hedge Standard and Poor's 500 index options because there is a very active market for the Standard and Poor's 500 futures that are traded on the Chicago Mercantile Exchange.

Since many of the stocks in the Standard and Poor's 500 index pay dividends, a time series of dividends for the index is necessary. The daily cash dividends for the index collected from the Standard and Poor's 500 information bulletin for the years 1992–94 can be used.¹ The present value of the dividends (until the option expires) is computed and subtracted from the current index level. For the risk-free rate, the continuously compounded Treasury bill rates (from the average of the bid and ask discounts reported in the *Wall Street Journal*) are interpolated to match the maturity of the option.

The raw intraday data set is sampled every Wednesday (or the next trading day if Wednesday is a holiday) between 2:30 P.M. and 3:15 P.M. central standard time to create the data set.² In particular, given a

particular Wednesday, an option must be traded on the following five trading days to be included in the sample. The study follows Dumas, Fleming, and Whaley (1998) in filtering the intraday data to create weekly data sets and use the midpoint of the bid-ask as the option price. As in Dumas, Fleming, and Whaley (1998), options with moneyness, $|K/F - 1|$ (K is the strike price and F is the forward price), less than or equal to 10 percent are included. In terms of maturity, options with time to maturity less than six days or greater than one hundred days are excluded.³

An option of a particular moneyness and maturity is represented only once in the sample on any particular day. In other words, although the same option may be quoted again in our time window (with the same or different index levels) on a given day, only the first record of that option is included in our sample for that day.

A transaction must satisfy the no-arbitrage relationship (Merton 1973) in that the call price must be greater than or equal to the spot price minus the present value of the remaining dividends and the discounted strike price. Similarly, the put price has to be greater than or equal to the present value of the remaining dividends plus the discounted strike price minus the spot price.

The entire data set consists of 7,404 records and observations spanning each trading day from January 5, 1994, to October 19, 1994.

1. Thanks to Jeff Fleming of Rice University for making the dividend series available.

2. Wednesdays are used as fewer holidays fall on Wednesdays.

3. See Dumas, Fleming, and Whaley (1998) for justification of the exclusionary criteria about moneyness and maturity.

20. If the number of options on a given day is too few, then there is a potential problem of overfitting in that more independent variables exist in the right-hand side but only a limited number of observations. However, such a problem can be partially mitigated by using a subset of the above regression (see Dumas, Fleming, and Whaley 1998).

21. One would want to test any options model in a very liquid options market so that prices are more reliable and do not reflect any liquidity premium.

22. The 1994 data were the latest full-year data available at the time of this writing.

23. The day t is usually a Wednesday. If Wednesday is a holiday, then the next trading day is chosen.

24. The mean absolute hedging error is the mean of the absolute values of the hedging errors. The conclusions do not change if a slightly different criterion is used, like root mean squared hedging error.

25. Prediction or out-of-sample valuation errors measure how well a given model values options based on the model parameters that were estimated in a previous time period.

Call option: Gives the owner of the option the right (but not the obligation) to buy the underlying asset at a fixed price (called the strike or exercise price). This right can be exercised at some fixed date in the future (European option) or at any time until the option matures (American option).

Put option: Gives the owner of the option the right (but not the obligation) to sell the underlying asset at a fixed price (called the strike or exercise price). This right can be exercised at some fixed date in the future (European option) or at any time until the option matures (American option).

Long position: In a security, implies that one has bought the security and currently owns it.

Short position: In a security, implies that one has sold a security that one does not own, but has only borrowed, with the hope of buying it back at a lower price in the future.

Implied volatility: The value of the volatility in the Black-Scholes-Merton formula that equates the model value of the option to its market price.

In-sample errors: Errors in fitting a model to data under a particular criterion function. For example, an options valuation model may have a few parameters or variables, the values of which are not observed directly. In such a case these parameters are estimated by minimizing a criterion function, such as the sum of squared differences between the model values and the market prices; this procedure is often called in-sample estimation. The differences between the model option values, evaluated at the estimates of the parameters, and the market option prices are called in-sample errors.

Out-of-sample errors: Measure the difference between the model option values and the market option prices on a sample of option prices that were observed at a later date than the sample on which the parameters of the model were estimated. In computing out-of-sample option values, the model parameters are fixed at the estimates obtained from the in-sample estimation.

implied but is computed from history of returns of the S&P 500 index), is quite competitive in terms of hedging across the entire sample of options.

Given the hedging results in the previous paragraph, which model would one choose among the four for constructing a hedge portfolio? The answer may very well depend on which option is to be hedged. The other panels of Table 4 show the mean absolute hedging errors of the four versions for three different classes of options: near-the-money call and put options and some relatively far-out-of-the-money put options. Most of these options are heavily traded in the Standard and Poor's 500 index options market.²⁶

The table shows that the differences in hedging errors among most of the versions are more clearly manifested in far-out-of-the-money put options. The ad hoc-1 version, in which the delta of an option is computed from its exact implied volatility, clearly dominates in terms of hedging out-of-the-money puts, irrespective of the maturity. For near-the-money options, the differences between the various versions are not that significant, especially if the portfolio is rebalanced on the next day. In fact, the least complex of all the versions, Black-Scholes-Merton-1, is quite competitive in terms of hedging near-the-money options.

Conclusion

Although the classic Black-Scholes-Merton paradigm of dynamic hedging is elegant from a theoretical perspective, it is often fraught with problems when it is implemented in the real world. Even if the Black-Scholes-Merton model were free of its known biases, the replicating/hedge portfolio of the model, which requires continuous trading, would rarely be able to match its target because trading can occur only at discrete intervals of time. Nevertheless, because of its simplicity and tractability, the model is widely used by options traders and investors. The basic delta-neutral hedge portfolio of the Black-Scholes-Merton model is also sometimes supplemented with other options to hedge a time-varying volatility (vega hedging). Although hedging a time-varying volatility is inconsistent with the Black-Scholes-Merton model, it can often prove useful in practice.

One would expect the presence of biases observed in the Black-Scholes-Merton model, such as the smile or smirk in implied volatilities, to result in further deterioration of the model's hedging performance. More advanced option pricing models (for example, random volatility models) that can account for some of the biases turn out to be useful mostly for deep out-of-the-money options but not

necessarily for near-the-money options. Ad hoc variations of the Black-Scholes-Merton model sometimes employed by options traders or investors to overcome the biases may also generate higher hedging errors than the very basic model despite the fact that ad hoc models often dominate the simple model in terms of matching observed option prices and predicting them. Although the simple Black-Scholes-Merton model can exhibit pricing biases, it is often competitive in terms of hedging because the pricing biases that it exhibits remain relatively stable through time.

Static hedging, an alternative to dynamic hedging, may seem promising because it is independent of any particular option pricing model. In particular, static hedging could prove useful for certain kinds of exotic options. However, static hedging requires hedging an option via other options so that the efficacy of static hedging depends on the liquidity of the options market, which often is not as liquid as the market on the underlying asset.

26. Far-out-of-the-money puts are those for which $K/F < 0.95$ where K is the strike price and F is the forward price for maturity τ —that is, $F(t) = S(t)\exp[r(\tau)t]$, where τ is the time to maturity of the option. Near-the-money options are those for which $|K/F - 1| \leq 0.01$.

REFERENCES

- BATES, DAVID. 1996. "Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutschemark Options." *Review of Financial Studies* 9 (Spring): 69–107.
- BLACK, FISCHER. 1976. "Studies of Stock Price Volatility Changes." In *Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economic Statistics Section*, 177–81. Alexandria, VA: American Statistical Association.
- BLACK, FISCHER, AND MYRON S. SCHOLES. 1973. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy* 81 (May/June): 637–54.
- CARR, PETER. 1998. "Randomization and the American Put." *Review of Financial Studies* 11 (Fall): 327–43.
- CHRISTIE, ANDREW A. 1982. "The Stochastic Behavior of Common Stock Variances: Value, Leverage, and Interest Rate Effects." *Journal of Financial Economics* 10 (December): 407–32.
- COX, JOHN C., AND MARK RUBINSTEIN. 1985. *Options Markets*. Englewood Cliffs, N.J.: Prentice-Hall.
- DERMAN, EMANUEL, AND IRAJ KANI. 1994. "Riding on the Smile." *Risk* 7 (February): 32–39.
- DUMAS, BERNARD, JEFF FLEMING, AND ROBERT WHALEY. 1998. "Implied Volatility Functions: Empirical Tests." *Journal of Finance* 53 (December): 2059–2106.
- DUPIRE, BRUNO. 1994. "Pricing with a Smile." *Risk* 7 (February): 18–20.
- HARRISON, J. MICHAEL, AND STANLEY R. PLISKA. 1981. "Martingales and Stochastic Integrals in the Theory of Continuous Trading." *Stochastic Processes and Their Applications* 11:215–60.
- HESTON, STEVEN L. 1993. "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options." *Review of Financial Studies* 6, no. 2:327–43.
- HESTON, STEVEN L., AND SAIKAT NANDI. Forthcoming. "A Closed-Form GARCH Option Valuation Model." *Review of Financial Studies*.
- HUANG, JING-ZHI, MARTI G. SUBRAHMANYAM, AND G. GEORGE YU. 1996. "Pricing and Hedging American Options: A Recursive Integration Method." *Review of Financial Studies* 9 (Spring): 277–300.
- HULL, JOHN C. 1997. *Options, Futures, and Other Derivatives*. 3d ed. Upper Saddle River, N.J.: Prentice-Hall.
- JACKWERTH, JENS, AND MARK RUBINSTEIN. 1996. "Recovering Probability Distributions from Option Prices." *Journal of Finance* 51 (December): 1611–52.
- JU, NENGJU. 1998. "Pricing an American Option by Approximating Its Early Exercise Boundary as a Multi-piece Exponential Function." *Review of Financial Studies* 11 (Fall): 627–46.
- MERTON, ROBERT C. 1973. "The Theory of Rational Option Pricing." *Bell Journal of Economics* 4 (Spring): 141–83.
- NANDI, SAIKAT. 1998. "How Important Is the Correlation between Returns and Volatility in a Stochastic Volatility Model? Empirical Evidence from Pricing and Hedging in the S&P 500 Index Options Market." *Journal of Banking and Finance* 22 (May): 589–610.
- RUBINSTEIN, MARK. 1994. "Implied Binomial Trees." *Journal of Finance* 49 (July): 771–818.
- VARIAN, HAL R. 1984. *Microeconomic Analysis*. New York: W.W. Norton and Company.