Anomalies in Option Pricing: The Black-Scholes Model Revisited

Peter Fortune

Professor of Economics, Tufts University and Senior Economist, Federal Reserve Bank of Boston. *The author is grateful to Lynn Browne, Lal Chugh, Dick Kopcke, Joe Peek and Bob Strong for their constructive comments on earlier drafts of this paper.*

This study is the third in a series of Federal Reserve Bank of Boston studies contributing to a broader understanding of derivative securities. The first (Fortune 1995) presented the rudiments of option pricing theory and addressed the equivalence between exchange-traded options and portfolios of underlying securities, making the point that plain vanilla options—and many other derivative securities—are really repackages of old instruments, not novel in themselves. That paper used the concept of portfolio insurance as an example of this equivalence. The second (Minehan and Simons 1995) summarized the presentations at "Managing Risk in the '90s: What Should You Be Asking about Derivatives?", an educational forum sponsored by the Boston Fed.

The present paper addresses the question of how well the best-known option pricing model—the Black-Scholes model—works. A full evaluation of the many option pricing models developed since their seminal paper in 1973 is beyond the scope of this paper. Rather, the goal is to acquaint a general audience with the key characteristics of a model that is still widely used, and to indicate the opportunities for improvement which might emerge from current research and which are undoubtedly the basis for the considerable current research on derivative securities. The hope is that this study will be useful to students of financial markets as well as to financial market practitioners, and that it will stimulate them to look into the more recent literature on the subject.

The paper is organized as follows. The next section briefly reviews the key features of the Black-Scholes model, identifying some of its most prominent assumptions and laying a foundation for the remainder of the paper. The second section employs recent data on almost one-half million options transactions to evaluate the Black-Scholes model. The third section discusses some of the reasons why the Black-Scholes model falls short and assesses some recent research designed to improve our ability to explain option prices. The paper ends with a brief summary. Those readers unfamiliar with the basics of stock options might refer to

Box 1: The Rudiments of Options on Common Stock

A **call option** gives the holder the right to acquire shares of a stock at the **exercise price**, also called the **strike price**, on or before a specific date, called the **expiration date**. The seller of the call, called the writer, is obligated to deliver the shares at the strike price if the option is exercised. A call option is said to be a **covered call** when the writer holds the shares that might have to be delivered upon exercise. A call option is **naked** when the writer does not own the underlying stock. Writing a covered call is roughly equivalent to writing a naked put option at the same strike price. Naked call options expose the option holder to the risk of nondelivery if the writer cannot buy the shares for delivery. Brokers typically require higher margins on naked calls.

A **put option** gives the holder the right to sell shares at the strike price on or before the exercise date. The writer of a put option is obligated to receive those shares and to deliver the required cash. A put option is **covered** if the writer has a short position in the underlying shares; otherwise, the put is **naked**. Writing a covered put is roughly equivalent to writing a naked call if the writer does not have a long position in the shares. Naked put options expose the option holder to the risk of loss if the writer does not have sufficient cash to pay for delivered shares.

The price paid for an option is called the **premium**. An option is said to be **in-the-money** if the holder would profit by exercising it; otherwise it is either **at-the-money** or **out-of-the-money**. Thus, a call option is in-the-money if the stock price exceeds the strike price, and it is out-of-the-money if the stock price is below the strike price. A put option is in-the-money if the stock price is below the strike price and out-of-the-money if the stock price exceeds the strike price. An option that remains out-of-the-money will not be exercised and will expire without any value.

An option is **European** if it can be exercised only on the expiration date. It is **American** if it can be exercised at any time on or before the expiration date. An **equity option** is an option on a specific firm's common stock. One equity option contract controls 100 shares of stock. When equity options are exercised, the resulting exchange is between cash and shares. All equity options traded on registered exchanges in the United States are American. An example of an equity option is the range of options on Intel's common stock. Traded at the American Stock Exchange, this option is available for several strike prices and expiration dates. For example, on February 2, 1996 there were transactions in the Intel call option expiring on February 16 with a strike price of $50 per share ($5,000 per contract). The premium at the **settlement** (close of trading) was $7 per share ($700 per contract). Because Intel's closing price on NASDAQ was $56.75 per share, this call option was in-the-money by $6.75 per share ($675 per contract).

A **stock index option** is an option on a stock index, and the resulting exchange is one of cash for cash. The holder of an exercised stock index option receives the difference between the S&P 500 at the time of exercise and the strike price, and the writer pays that amount. Each index futures contract is for $100 times the value of the index. All stock index options traded in the United States are American with one significant exception: The S&P 500-stock index option is European. Denoted as **SPX**, the S&P 500 index option is traded on the Chicago Board of Trade's Option Exchange (CBOE). On February 2, 1996, the CBOE's SPX index option was traded for expiration dates from February 1996 through December 1997. For each expiration date there were options at a range of strike prices. On February 2, 1996, when the S&P 500 closed at 635.84, the settlement premium on the February SPX call option with a strike price of 640 ($64,000 per contract) was 5.00 ($500 per contract). The option was out-of-the-money, because if it were immediately exercised, the holder would receive $63,584, for which he would pay $64,000.

All traded options expire on the third Friday of their exercise month. Option contracts are not written directly between the buyer and seller. Instead, each party makes a contract with a **clearing house**. In the United States the Option Clearing Corporation is the major clearing house. The primary function of the clearing house is to eliminate counterparty risk as a significant consideration. That is, the option holder need not fear that the writer will not honor the option, because the clearing house will honor it. If the holder of an option chooses to exercise it, the clearing house will randomly select a writer of the same type of option to make delivery.

Notation

C: the premium on a call option
P: the premium on a put option
S: the price of the underlying security
X: the option's strike price
r: the riskless rate of interest
σ: the option's volatility
T: the option's expiration date
t: the current date

Fortune (1995). Box 1 reviews briefly the fundamental language of options and explains the notation used in the paper.

I. The Black-Scholes Model

In 1973, Myron Scholes and the late Fischer Black published their seminal paper on option pricing (Black and Scholes 1973). The Black-Scholes model revolutionized financial economics in several ways. First, it contributed to our understanding of a wide range of contracts with option-like features. For example, the call feature in corporate and municipal bonds is clearly an option, as is the refinancing privilege in mortgages. Second, it allowed us to revise our understanding of traditional financial instruments. For example, because shareholders can turn the company over to creditors if it has negative net worth, corporate debt can be viewed as a put option bought by the shareholders from creditors.

The Black-Scholes model explains the prices on European options, which cannot be exercised before the expiration date. Box 2 summarizes the Black-Scholes model for pricing a European call option on which dividends are paid continuously at a constant

Box 2: The Black-Scholes Option Pricing Model with Continuous Dividends

Following Merton (1973), we consider a share of common stock that pays a continuous dividend at a constant yield of q at each moment, and a call option that expires at time T. The current price of a share, at time t, is denoted as S_t. This price can be interpreted as the sum of two components. The first component is the present value of the dividends to be paid over the period up to time T, which is the expiration date of a call option on the stock. The second component is the value that is "at risk." Because payment of dividends reduces the value of the stock at the rate q, the stock price at time T is reduced by the factor $e^{-q(T-t)}$, so the present value "at-risk" is $S_t e^{-q(T-t)}$.

Denoting the "at-risk" component as S^*, the Black-Scholes model assumes that S^* evolves over time as a diffusion process, which can be written as

$$dS^*/S^* = \mu dt + \sigma dz \qquad \text{(B2.1)}$$

in which μ, called the "drift," is the expected instantaneous rate of change in S^*, and σ, called the "volatility," is the standard deviation of the instantaneous rate of change in S^*. The term dz, called a Wiener variable, is a normally distributed random variable with a mean of zero and a standard deviation of $\sqrt{dt}$. Thus, the rate of change in S^* vibrates randomly around the drift. If we convert this to a statement about the value of S^*, we find that S^* will be log-normally distributed, that is, the logarithm of S^* will be normally distributed.

Now consider a European call option on that stock which expires in $(T - t)$ days. The Black-Scholes model describes the equilibrium price, or premium, on an option as a function of the risky component of the stock price $(S_t e^{-q(T-t)})$, the present value of the option's strike price $(Xe^{-r(T-t)})$, the riskless rate of interest (r), the dividend-yield on the stock (q), the time remaining until the option expires $(T - t)$, and the "volatility" of the return on the underlying security (σ). The volatility is defined as the standard deviation of the rate of change in the stock's price.

Recalling that $S_t^* = S_t e^{-q(T-t)}$, the Black-Scholes relationship is

$$C_t = S_t^* N(d_1) - Xe^{-r(T-t)} N(d_2) \qquad \text{(B2.2)}$$

where

$$d_1 = [\ln(S/X) + (r - q + \tfrac{1}{2}\sigma^2)(T - t)]/\sigma\sqrt{(T - t)}$$

$$d_2 = d_1 - \sigma\sqrt{(T - t)}$$

In this formula N(d) is the probability that a standard normal random variable is less than d. $N(d_1)$ and $N(d_2)$, both positive but less than one, represent the number of shares and the amount of debt in a portfolio that exactly replicate the price of the option. Thus, a call option on one share is exactly equivalent to buying $N(d_1)$ shares of the stock and selling $N(d_2)$ units of a bond with present value $Xe^{-r(T-t)}$. For example, if $N(d_1) = 0.5$ and $N(d_2) = 0.4$, the call option is exactly equivalent to one-half share of the stock plus borrowing 40 percent of the present value of the strike price; this is the option's "replicating portfolio" and a position consisting of one call option, shorting $N(d_1)$ shares, and purchasing $Xe^{-r(T-t)} N(d_2)$ of bonds creates a perfect hedge, exposing the holder to no price risk.

rate. A crucial feature of the model is that the call option is equivalent to a portfolio constructed from the underlying stock and bonds. The "option-replicating portfolio" consists of a fractional share of the stock combined with borrowing a specific amount at the riskless rate of interest. This equivalence, developed more fully in Fortune (1995), creates price relationships which are maintained by the arbitrage of informed traders. The Black-Scholes option pricing model is derived by identifying an option-replicating portfolio, then equating the option's premium with the value of that portfolio.

An essential assumption of this pricing model is that investors arbitrage away any profits created by gaps in asset pricing. For example, if the call is trading "rich," investors will write calls and buy the replicating portfolio, thereby forcing the prices back into line. If the option is trading low, traders will buy the option and short the option-replicating portfolio (that is, sell stocks and buy bonds in the correct proportions). By doing so, traders take advantage of riskless opportunities to make profits, and in so doing they force option, stock, and bond prices to conform to an equilibrium relationship.

The Black-Scholes model revolutionized financial economics in several ways, contributing to our understanding of a wide range of contracts with option-like features.

Arbitrage allows European puts to be priced using *put-call parity*. Consider purchasing one call that expires at time T and lending the present value of the strike price at the riskless rate of interest. The cost is $C_t + Xe^{-r(T-t)}$. (See Box 1 for notation: C is the call premium, X is the call's strike price, r is the riskless interest rate, T is the call's expiration date, and t is the current date.) At the option's expiration the position is worth the highest of the stock price (S_T) or the strike price, a value denoted as $max(S_T, X)$. Now consider another investment, purchasing one put with the same strike price as the call, plus buying the fraction $e^{-q(T-t)}$ of one share of the stock. Denoting the put premium by P and the stock price by S, then the cost of this is $P_t + e^{-q(T-t)}S_t$, and, at time T, the value at this position is also $max(S_T, X)$.[1] Because both positions have the same terminal value, arbitrage will force them to have the same initial value.

Suppose that $C_t + Xe^{-r(T-t)} > P_t + e^{-q(T-t)}S_t$, for example. In this case, the cost of the first position exceeds the cost of the second, but both must be worth the same at the option's expiration. The first position is overpriced relative to the second, and shrewd investors will go short the first and long the second; that is, they will write calls and sell bonds (borrow), while simultaneously buying both puts and the underlying stock. The result will be that, in equilibrium, equality will prevail and $C_t + Xe^{-r(T-t)} = P_t + e^{-q(T-t)}S_t$. Thus, arbitrage will force a parity between premiums of put and call options.

Using this put-call parity, it can be shown that the premium for a European put option paying a continuous dividend at q percent of the stock price is:

$$P_t = -e^{-q(T-t)}S_tN(-d_1) + Xe^{-r(T-t)}N(-d_2)$$

where d_1 and d_2 are defined as in Box 2.

The importance of arbitrage in the pricing of options is clear. However, many option pricing models can be derived from the assumption of complete arbitrage. Each would differ according to the probability distribution of the price of the underlying asset. What makes the Black-Scholes model unique is that it assumes that stock prices are log-normally distributed, that is, that the logarithm of the stock price is normally distributed. This is often expressed in a "diffusion model" (see Box 2) in which the (instantaneous) rate of change in the stock price is the sum of two parts, a "drift," defined as the difference between the expected rate of change in the stock price and the dividend yield, and "noise," defined as a random variable with zero mean and constant variance. The variance of the noise is called the "volatility" of the stock's rate of price change. Thus, the rate of change in a stock price vibrates randomly around its expected value in a fashion sometimes called "white noise."

The Black-Scholes models of put and call option pricing apply directly to *European* options as long as a continuous dividend is paid at a constant rate. If no

[1] Consider the call *cum* bond position, purchased for $C_t + Xe^{-r(T-t)}$. If, at expiration, $S_T \leq X$, the call will expire without value and the position will be worth the accumulated value of the bond, or X. However, if, at expiration, the call is in-the-money (that is, $S_T > X$), it will be exercised and the holder will receive $S_T - X$. When added to the value of the bond at time T, the position is worth $S_T - X + X = S_T$. Thus, the call *cum* bond position is worth the highest of S_T or X, a value denoted by $max(S_T, X)$.

Figure 1

Call Premium and Stock Price

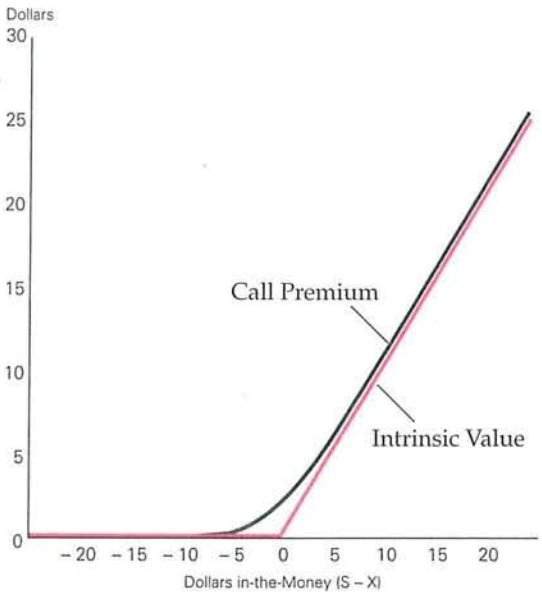

dividends are paid, the models also apply to *American* call options, which can be exercised at any time. In this case, it can be shown that there is no incentive for early exercise, hence the American call option must trade like its European counterpart. However, the Black-Scholes model does not hold for American put options, because these might be exercised early, nor does it apply to any American option (put or call) when a dividend is paid.[2] Our empirical analysis will sidestep those problems by focusing on European-style options, which cannot be exercised early.

A call option's intrinsic value is defined as $max(S - X,0)$, that is, the largest of $S - X$ or zero; a put option's intrinsic value is $max(X - S,0)$. When the stock price (S) exceeds a call option's strike price (X), or falls short of a put option's strike price, the option has a positive intrinsic value because if it could be immediately exercised, the holder would receive a gain of $S - X$ for a call, or $X - S$ for a put. However, if $S < X$, the holder of a call will not exercise the option and it has no intrinsic value; if $X > S$ this will be true for a put.

The intrinsic value of a call is the kinked line in Figure 1 (a put's intrinsic value, not shown, would have the opposite kink). When the stock price exceeds the strike price, the call option is said to be in-the-money. It is out-of-the-money when the stock price is below the strike price. Thus, the kinked line, or intrinsic value, is the income from immediately exercising the option: When the option is out-of-the-money, its intrinsic value is zero, and when it is in the money, the intrinsic value is the amount by which S exceeds X.

Convexity, the Call Premium, and the Greek Chorus

The premium, or price paid for the option, is shown by the curved line in Figure 1. This curvature, or "convexity," is a key characteristic of the premium on a call option. Figure 1 shows the relationship between a call option's premium and the underlying stock price for a hypothetical option having a 60-day term, a strike price of \$50, and a volatility of 20 percent. A 5 percent riskless interest rate is assumed. The call premium has an upward-sloping relationship with the stock price, and the slope rises as the stock price rises. This means that the sensitivity of the call premium to changes in the stock price is not constant and that the option-replicating portfolio changes with the stock price.

The convexity of option premiums gives rise to a number of technical concepts which describe the response of the premium to changes in the variables and parameters of the model. For example, the relationship between the premium and the stock price is captured by the option's Delta (Δ) and its Gamma (Γ). Defined as the slope of the premium at each stock price, the Delta tells the trader how sensitive the option price is to a change in the stock price.[3] It also tells the trader the value of the hedging ratio.[4] For each share of stock held, a perfect hedge requires writing $1/\Delta_c$ call options or buying $1/\Delta_p$ puts. Figure 2 shows the Delta for our hypothetical call option as a function of the stock price. As S increases, the value of Delta rises until it reaches its maximum at a stock

---

[2] If a dividend is paid, an American call option might be exercised early to capture the dividend. American puts might be exercised early regardless of a dividend payment if they are deep-in-the-money. Thus, American options might be priced differently from European options.

[3] Delta is defined as $\Delta_c = \partial C/\partial S$ for a call and $\Delta_p = \partial P/\partial S$ for a put.

[4] The hedging ratio is the number of options that must be written or bought to insulate the investor from the effects of a change in the price of a share of the underlying stock. Thus, if $\Delta_c = 0.33$, the hedging ratio using calls is -3, that is, calls on 300 shares (3 contracts) must be written to protect 100 shares against a change in the stock price.

Figure 2

The Greek Chorus and the Stock Price

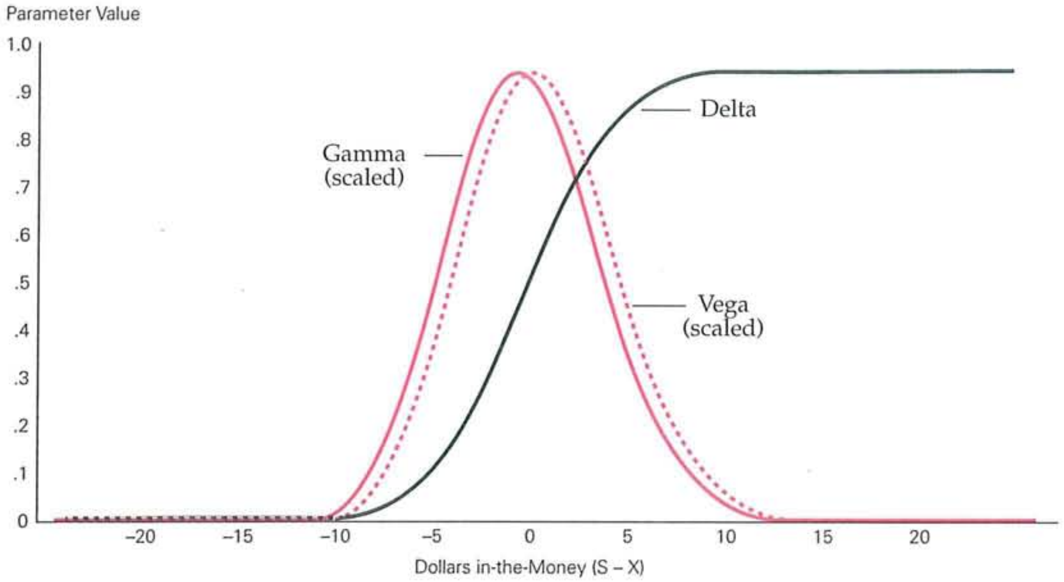

price of about \$60, or \$10 in-the-money. After that point, the option premium and the stock price have a 1:1 relationship. The increasing Delta also means that the hedging ratio falls as the stock price rises. At higher stock prices, fewer call options need to be written to insulate the investor from changes in the stock price.

The Gamma is the change in the Delta when the stock price changes.[5] Gamma is positive for calls and negative for puts. The Gamma tells the trader how much the hedging ratio changes if the stock price changes. If Gamma is zero, Delta would be independent of S and changes in S would not require adjustment of the number of calls required to hedge against further changes in S. The greater is Gamma, the more "out-of-line" a hedge becomes when the stock price changes, and the more frequently the trader must adjust the hedge.

Figure 2 shows the value of Gamma as a function of the amount by which our hypothetical call option is in-the-money.[6] Gamma is almost zero for deep-in-the-money and deep-out-of-the-money options, but it reaches a peak for near-the-money options. In short, traders holding near-the-money options will have to adjust their hedges frequently and sizably as the stock price vibrates. If traders want to go on long vacations without changing their hedges, they should focus on far-away-from-the-money options, which have near-zero Gammas.

A third member of the Greek chorus is the option's Lambda, denoted by Λ, also called Vega.[7] Vega measures the sensitivity of the call premium to changes in volatility. The Vega is the same for calls and puts having the same strike price and expiration

[5] $\Gamma_c = \partial\Delta_c/\partial S = \partial^2 C/\partial S^2$ for a call and $\Gamma_p = \partial\Delta_p/\partial S = \partial^2 P/\partial S^2$ for a put.

[6] Because the actual values of Delta, Gamma, and Vega are very different, some scaling is necessary to put them on the same figure. We have scaled by dividing actual values by the maximum value. Thus, each curve in Figure 2 shows the associated parameter relative to its peak value, with the peak set to 1. Note that Delta is already scaled since its maximum is 1.

[7] Vega is not a Greek letter, but it serves as a useful mnemonic for the sensitivity of the premium to changes in Volatility. $\Lambda_c = \partial C/\partial\sigma$ for a call and $\Lambda_p = \partial P/\partial\sigma$ for a put, where σ is the volatility of the daily return on the stock.

date. As Figure 2 shows, a call option's Vega conforms closely to the pattern of its Gamma, peaking for near-the-money options and falling to zero for deep-out or deep-in options. Thus, near-the-money options appear to be most sensitive to changes in volatility.

Because an option's premium is directly related to its volatility—the higher the volatility, the greater the chance of it being deep-in-the-money at expiration—any propositions about an option's price can be translated into statements about the option's volatility, and vice versa. For example, other things equal, a high volatility is synonymous with a high option premium for both puts and calls. Thus, in many contexts we can use volatility and premium interchangeably. We will use this result below when we address an option's *implied volatility*.

Other Greeks are present in the Black-Scholes pantheon, though they are lesser gods. The option's Rho (ρ) is the sensitivity of the call premium to changes in the riskless interest rate.[8] Rho is always positive for a call (negative for a put) because a rise in the interest rate reduces the present value of the strike price paid (or received) at expiration if the option is exercised. The option's Theta (θ) measures the change in the premium as the term shortens by one time unit.[9] Theta is always negative because an option is less valuable the shorter the time remaining.

The Black-Scholes Assumptions

The assumptions underlying the Black-Scholes model are few, but strong. They are:

- *Arbitrage*: Traders can, and will, eliminate any arbitrage profits by simultaneously buying (or writing) options and writing (or buying) the option-replicating portfolio whenever profitable opportunities appear.
- *Continuous Trading*: Trading in both the option and the underlying security is continuous in time, that is, transactions can occur simultaneously in related markets at any instant.
- *Leverage*: Traders can borrow or lend in unlimited amounts at the riskless rate of interest.
- *Homogeneity*: Traders agree on the values of the relevant parameters, for example, on the riskless rate of interest and on the volatility of the returns on the underlying security.
- *Distribution*: The price of the underlying security is log-normally distributed with statistically independent price changes, and with constant mean and constant variance.
- *Continuous Prices*: No discontinuous jumps occur in the price of the underlying security.
- *Transactions Costs*: The cost of engaging in arbitrage is negligibly small.

The *arbitrage* assumption, a fundamental proposition in economics, has been discussed above. The *continuous trading* assumption ensures that at all times traders can establish hedges by simultaneously trading in options and in the underlying portfolio. This is important because the Black-Scholes model derives its power from the assumption that at any instant, arbitrage will force an option's premium to be equal to the value of the replicating portfolio. This cannot be done if trading occurs in one market while trading in related markets is barred or delayed. For example, during a halt in trading of the underlying security one would not expect option premiums to conform to the Black-Scholes model. This would also be true if the underlying security were inactively traded, so that the trader had "stale" information on its price when contemplating an options transaction.

If traders want to go away on long vacations without changing their hedges, they should focus on far-from-the-money options, which have near-zero Gammas.

The *leverage* assumption allows the riskless interest rate to be used in options pricing without reference to a trader's financial position, that is, to whether and how much he is borrowing or lending. Clearly this is an assumption adopted for convenience and is not strictly true. However, it is not clear how one would proceed if the rate on loans was related to traders' financial choices. This assumption is common to finance theory: For example, it is one of the assumptions of the Capital Asset Pricing Model. Furthermore, while private traders have credit risk, important players in the option markets, such as nonfinancial corporations and major financial institutions, have very low credit risk over the lifetime of most options (a year or less), suggesting that departures from this assumption might not be very important.

The *homogeneity* assumption, that traders share

[8] $\rho_c = \partial C/\partial r$ and $\rho_p = \partial P/\partial r$.
[9] $\theta_c = \partial C/\partial t$ and $\theta_p = \partial C/\partial t$.

the same probability beliefs and opportunities, flies in the face of common sense. Clearly, traders differ in their judgments of such important things as the volatility of an asset's future returns, and they also differ in their time horizons, some thinking in hours, others in days, and still others in weeks, months, or years. Indeed, much of the actual trading that occurs must be due to differences in these judgments, for otherwise there would be no disagreements with "the market" and financial markets would be pretty dull and uninteresting.

The *distribution* assumption is that stock prices are generated by a specific statistical process, called a diffusion process, which leads to a normal distribution of the *logarithm* of the stock's price. Furthermore, the *continuous price* assumption means that any changes in prices that are observed reflect only different draws from the same underlying log-normal distribution, not a change in the underlying probability distribution itself.

II. Tests of the Black-Scholes Model

Assessments of a model's validity can be done in two ways. First, the model's predictions can be confronted with historical data to determine whether the predictions are accurate, at least within some statistical standard of confidence. Second, the assumptions made in developing the model can be assessed to determine if they are consistent with observed behavior or historical data.

A long tradition in economics focuses on the first type of tests, arguing that "the proof is in the pudding." It is argued that any theory requires assumptions that might be judged "unrealistic," and that if we focus on the assumptions, we can end up with no foundations for deriving the generalizations that make theories useful. The only proper test of a theory lies in its predictive ability: The theory that consistently predicts best is the best theory, regardless of the assumptions required to generate the theory.

Tests based on assumptions are justified by the principle of "garbage in-garbage out." This approach argues that no theory derived from invalid assumptions can be valid. Even if it appears to have predictive abilities, those can slip away quickly when changes in the environment make the invalid assumptions more pivotal.

Our analysis takes an agnostic position on this methodological debate, looking at both predictions and assumptions of the Black-Scholes model.

The Data

The data used in this study are from the Chicago Board Options Exchange's Market Data Retrieval System. The MDR reports the number of contracts traded, the time of the transaction, the premium paid, the characteristics of the option (put or call, expiration date, strike price), and the price of the underlying stock at its last trade. This information is available for each option listed on the CBOE, providing as close to a real-time record of transactions as can be found. While our analysis uses only records of actual transactions, the MDR also reports the same information for every request of a quote. Quote records differ from the transaction records only in that they show both the bid and asked premiums and have a zero number of contracts traded.

Assessments of a model's validity can be done in two ways. The model's predictions can be confronted with historical data, or the assumptions made in developing the model can be assessed.

The data used are for the 1992–94 period. We selected the MDR data for the S&P 500-stock index (SPX) for several reasons. First, the SPX options contract is the only European-style stock index option traded on the CBOE. All options on individual stocks and on other indices (for example, the S&P 100 index, the Major Market Index, the NASDAQ 100 index) are American options for which the Black-Scholes model would not apply. The ability to focus on a European-style option has several advantages. By allowing us to ignore the potential influence of early exercise, a possibility that significantly affects the premiums on American options on dividend-paying stocks as well as the premiums on deep-in-the-money American put options, we can focus on options for which the Black-Scholes model was designed. In addition, our interest is not in individual stocks and their options, but in the predictive power of the Black-Scholes option pricing model. Thus, an index option allows us to make broader generalizations about model performance than would a select set of equity options. Finally, the S&P 500 index options trade in a very active market, while options on many individual stocks and on some other indices are thinly traded.

The full MDR data set for the SPX over the roughly 758 trading days in the 1992–94 period consisted of more than 100 million records. In order to bring this down to a manageable size, we eliminated all records that were requests for quotes, selecting only records reflecting actual transactions. Some of these transaction records were cancellations of previous trades, for example, trades made in error. If a trade was canceled, we included the records of the original transaction because they represented market conditions at the time of the trade, and because there is no way to determine precisely which transaction was being canceled. We eliminated cancellations because they record the S&P 500 at the time of the cancellation, not the time of the original trade. Thus, cancellation records will contain stale prices.

The Black-Scholes model assumes that investors know the volatility of the rate of return on the underlying asset, and that an option's implied volatility should differ from the true volatility only because of random events.

This screening created a data set with over 726,000 records. In order to complete the data required for each transaction, the bond-equivalent yield (average of bid and asked prices) on the Treasury bill with maturity closest to the expiration date of the option was used as a riskless interest rate. These data were available for 180-day terms or less, so we excluded options with a term longer than 180 days, leaving over 486,000 usable records having both CBOE and Treasury bill data. For each of these, we assigned a dividend yield based on the S&P 500 dividend yield in the month of the option trade.

Because each record shows the actual S&P 500 at almost the same time as the option transaction, the MDR provides an excellent basis for estimating the theoretically correct option premium and evaluating its relationship to actual option premiums. There are, however, some minor problems with interpreting the MDR data as providing a trader's-eye view of option pricing. The transaction data are not entered into the CBOE computer at the exact moment of the trade. Instead, a ticket is filled out and then entered into the computer, and it is only at that time that the actual level of the S&P 500 is recorded. In short, the S&P 500 entries necessarily lag behind the option premium entries, so if the S&P 500 is rising (falling) rapidly, the reported value of the SPX will be above (below) the true value known to traders at the time of the transaction.

Test 1: An Implied Volatility Test

A key variable in the Black-Scholes model is the volatility of returns on the underlying asset, the SPX in our case. Investors are assumed to know the true standard deviation of the rate of return over the term of the option, and this information is embedded in the option premium. While the true volatility is an unobservable variable, the market's estimate of it can be inferred from option premiums. The Black-Scholes model assumes that this "implied volatility" is an optimal forecast of the volatility in SPX returns observed over the term of the option.

The calculation of an option's implied volatility is reasonably straightforward. Six variables are needed to compute the predicted premium on a call or put option using the Black-Scholes model. Five of these can be objectively measured within reasonable tolerance levels: the stock price (S), the strike price (X), the remaining life of the option ($T - t$), the riskless rate of interest over the remaining life of the option (r), typically measured by the rate of interest on U.S. Treasury securities that mature on the option's expiration date, and the dividend yield (q). The sixth variable, the "volatility" of the return on the stock price, denoted by σ, is unobservable and must be estimated using numerical methods. Using reasonable values of all the known variables, the implied volatility of an option can be computed as the value of σ that makes the predicted Black-Scholes premium exactly equal to the actual premium. An example of the computation of the implied volatility on an option is shown in Box 3.

The Black-Scholes model assumes that investors know the volatility of the rate of return on the underlying asset, and that this volatility is measured by the (population) standard deviation. If so, an option's implied volatility should differ from the true volatility only because of random events. While these discrep-

Box 3: Computing Implied Volatility

At 8:55:02 a.m. on December 30, 1994 a transaction was recorded for 100 contracts on SPX calls having a strike price of 460 and expiring on March 17, 1995. The call premium was 12.75 and the value of the S&P 500 was recorded as 461.93 at the time of the trade. The option had a term of 77 days, the annual interest rate on Treasury bills that expired closest to that date was 5.56 percent (0.00015233 per day), and the dividend-yield prevailing at the time was 3.01 percent (0.00008247 per day).

The Black-Scholes model for this call option (see Box 2) allows us to compute the implied volatility. For this option we have $S = 461.93$, $X = 460$, $C = 12.75$, $T - t = 77$, $r = 0.00015233$ and $q = 0.00008247$. Denoting the daily volatility as σ, the model in Box 2 gives

$$d_1 = [ln(461.93/460) + (0.00015233 - 0.00008247 + \tfrac{1}{2}\sigma^2)(77)]/(\sigma\sqrt{77}) \qquad \text{(B3.1)}$$

$$= (0.0011156 + 4.387482\ \sigma^2)/\sigma$$

$$d_2 = d1 - 8.77496\sigma$$

For each possible value of σ the values of $N(d_1)$ and $N(d_2)$ can be computed from the standard normal distribution function. These values can then be fed into the implicit equation

$$C - [Se^{-q(T-t)}N(d_1) - Xe^{-r(T-t)}N(d_2)] = 0 \qquad \text{(B3.2)}$$

where C is the actual premium and the term in brackets is the Black-Scholes theoretical premium. Different values of σ can be tried until the theoretical and actual premiums are equal. The solution for the daily volatility is the value of σ that solves equation (B3.2), hence it is the daily standard deviation that makes the Black-Scholes model explain the observed premium. To convert this to a percentage value at annual rates we multiply σ by $100\sqrt{253}$ (following the convention of using a 253-day trading year). For our particular option, the implied volatility is 10.3 percent.

Computations for the 486,000 options transactions in our sample were done using the OPTMUM module for the statistical and econometrics program GAUSS.

ancies might occur, they should be very short-lived and random: Informed investors will observe the discrepancy and engage in arbitrage, which quickly returns things to their normal relationships.

Figure 3 reports two measures of the volatility in the rate of return on the S&P 500 index for each trading day in the 1992–94 period.[10] The "actual" volatility is the ex post standard deviation of the daily change in the logarithm of the S&P 500 over a 60-day horizon, converted to a percentage at an annual rate. For example, for January 5, 1993 the standard deviation of the daily change in *ln*S&P500 was computed for the next 60 calendar days; this became the actual volatility for that day. Note that the actual volatility is the realization of one outcome from the entire probability distribution of the standard deviation of the rate of return. While no single realization will be equal to the "true" volatility, the actual volatility should equal the true volatility, "on average."

The second measure of volatility is the implied volatility. This was constructed as follows, using the data described above. For each trading day, the implied volatility on call options meeting two criteria was computed. The criteria were that the option had 45 to 75 calendar days to expiration (the average was 61 days) and that it be near the money (defined as a spread between S&P 500 and strike price no more than 2.5 percent of the S&P 500). The first criterion was adopted to match the term of the implied volatility with the 60-day term of the actual volatility. The second criterion was chosen because, as we shall see later, near-the-money options are most likely to conform to Black-Scholes predictions.

The Black-Scholes model assumes that an option's implied volatility is an optimal forecast of the volatility in SPX returns observed over the term of the option. Figure 3 does not provide visual support for the idea that implied volatilities deviate randomly from actual volatility, a characteristic of optimal forecasting. While the two volatility measures appear to have roughly the same average, extended periods of significant differences are seen. For example, in the last half of 1992 the implied volatility remained well above the actual volatility, and after the two came together in the first half of 1993, they once again diverged for an extended period. It is clear from this visual record that implied volatility does not track actual volatility

[10] This conversion is done by multiplying the daily value of σ by $\sqrt{253}$ to bring it to an annual rate based on 253 trading days per year. The result is then multiplied by 100 to convert from fractions to percentages.

Figure 3

Volatility of 60-Day S&P 500 Returns

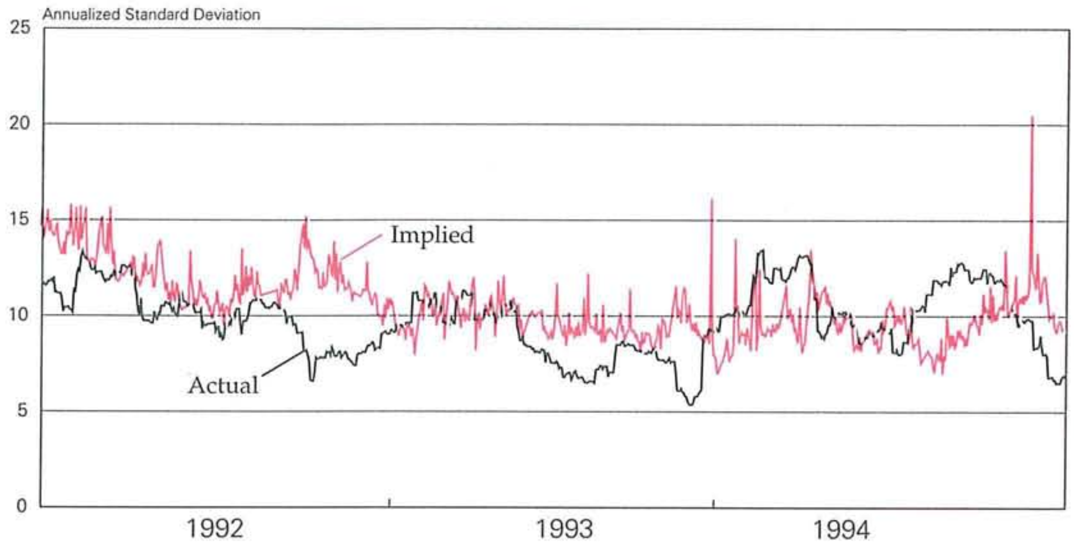

Source: Author's calculations.

well. However, this does not mean that implied volatility provides an inferior forecast of actual volatility: It could be that implied volatility satisfies all the scientific requirements of a good forecast in the sense that no other forecasts of actual volatility are better.

In order to pursue the question of the informational content of implied volatility, several simple tests of the hypothesis that implied volatility is an optimal forecast of actual volatility can be applied. One characteristic of an optimal forecast is that the forecast should be unbiased, that is, the forecast error (actual volatility less implied volatility) should have a zero mean. The average forecast error for the data shown in Figure 3 is −0.7283, with a t-statistic of −8.22. This indicates that implied volatility is a biased forecast of actual volatility.

A second characteristic of an optimal forecast is that the forecast error should not depend on any information available at the time the forecast is made. If information were available that would improve the forecast, the forecaster should have already included it in making his forecast. Any remaining forecasting errors should be random and uncorrelated with information available before the day of the forecast. To implement this "residual information test," the forecast error was regressed on the lagged values of the S&P 500 in the three days prior to the forecast.[11] The F-statistic for the significance of the regression coefficients was 4.20, with a significance level of 0.2 percent. This is strong evidence of a statistically significant violation of the residual information test.

The conclusion that implied volatility is a poor forecast of actual volatility has been reached in several other studies using different methods and data. For example, Canina and Figlewski (1993), using data for the S&P 100 in the years 1983 to 1987, found that implied volatility had almost no informational content as a prediction of actual volatility. However, a recent

[11] Because the forecast errors are for overlapping periods, they must be serially correlated. This calls for an estimation method that corrects for serial correlation. The method used was Hannan Efficient regression, in which serial correlation in residuals is corrected using spectral analysis. Over the 758-day period between 1992 and 1994, on 67 days no forecast error could be measured because no at-the-money call options with 45 to 75 days remaining were traded. The average forecast error over the 691 days with measured forecast errors was substituted for the missing values on those days. The independent variables were a dummy variable for missing forecast error and the lagged values of the SPX over the previous three days.

review of the literature on implied volatility (Mayhew 1995) mentions a number of papers that give more support for the forecasting ability of implied volatility.

Test 2: The Smile Test

One of the predictions of the Black-Scholes model is that at any moment all SPX options that differ only in the strike price (having the same term to expiration) should have the same implied volatility. For example, suppose that at 10:15 a.m. on November 3, transactions occur in several SPX call options that differ only in the strike price. Because each of the options is for the same interval of time, the value of volatility embedded in the option premiums should be the same. This is a natural consequence of the fact that the variability in the S&P 500's return over any future period is independent of the strike price of an SPX option.

One approach to testing this is to calculate the implied volatilities on a set of options identical in all respects except the strike price. If the Black-Scholes model is valid, the implied volatilities should all be the same (with some slippage for sampling errors). Thus, if a group of options all have a "true" volatility of, say, 12 percent, we should find that the implied volatilities differ from the true level only because of random errors. Possible reasons for these errors are temporary deviations of premiums from equilibrium levels, or a lag in the reporting of the trade so that the value of the SPX at the time stamp is not the value at the time of the trade, or that two options might have the same time stamp but one was delayed more than the other in getting into the computer.

This means that a graph of the implied volatilities against any economic variable should show a flat line. In particular, no relationship should exist between the implied volatilities and the strike price or, equivalently, the amount by which each option is "in-the-money." However, it is widely believed that a "smile" is present in option prices, that is, options far out of the money or far in the money have higher implied volatilities than near-the-money options. Stated differently, deep-out and far-in options trade "rich" (overpriced) relative to near-the-money options.

If true, this would make a graph of the implied volatilities against the value by which the option is in-the-money look like a smile: high implied volatilities at the extremes and lower volatilities in the middle. In order to test this hypothesis, our MDR data were screened for each day to identify any options that have the same characteristics but different strike prices. If 10 or more of these "identical" options were found, the average implied volatility for the group was computed and the deviation of each option's implied volatility from its group average, the Volatility Spread, was computed. For each of these options, the amount by which it is in-the-money was computed, creating a variable called ITM (an acronym for in-the-money). ITM is the amount by which an option is in-the-money. It is negative when the option is out-of-the-money. ITM is measured relative to the S&P 500 index level, so it is expressed as a percentage of the S&P 500.

The Volatility Spread was then regressed against a fifth-order polynomial equation in ITM. This allows for a variety of shapes of the relationship between the two variables, ranging from a flat line if Black-Scholes is valid (that is, if all coefficients are zero), through a wavy line with four peaks and troughs. The Black-Scholes prediction that each coefficient in the polynomial regression is zero, leading to a flat line, can be tested by the F-statistic for the regression. The results are reported in Table 1, which shows the F-statistic for the hypothesis that all coefficients of the fifth-degree polynomial are jointly zero. Also reported is the proportion of the variation in the Volatility Spreads, which is explained by variations in ITM (R^2). The results strongly reject the Black-Scholes model. The F-statistics are extremely high, indicating virtually no chance that the value of ITM is irrelevant to the explanation of implied volatilities. The values of R^2 are also high, indicating that ITM explains about 40 to 60 percent of the variation in the Volatility Spread.

Table 1
Testing for a Smile[a]

Year	Type	Number of Trades	F-Stat[b]	DF	R^2
1992	Call	43,449	5,561	5/43,443	.39
	Put	65,267	14,890	5/65,261	.53
1993	Call	59,269	8,758	5/59,263	.43
	Put	88,501	24,934	5/88,495	.58
1994	Call	82,828	9,530	5/82,822	.37
	Put	137,640	47,528	5/137,634	.63

[a]Option transactions with volatility spreads below the first percentile or greater than the 99th percentile were excluded in order to eliminate the influence of outliers.

[b]The F-statistics are for regressions of the deviation of a volatility from its group mean on a fifth-degree polynomial in ITM, the amount by which the option is in-the-money divided by the SPX level.

Figure 4

The Smile in Call Option Volatility

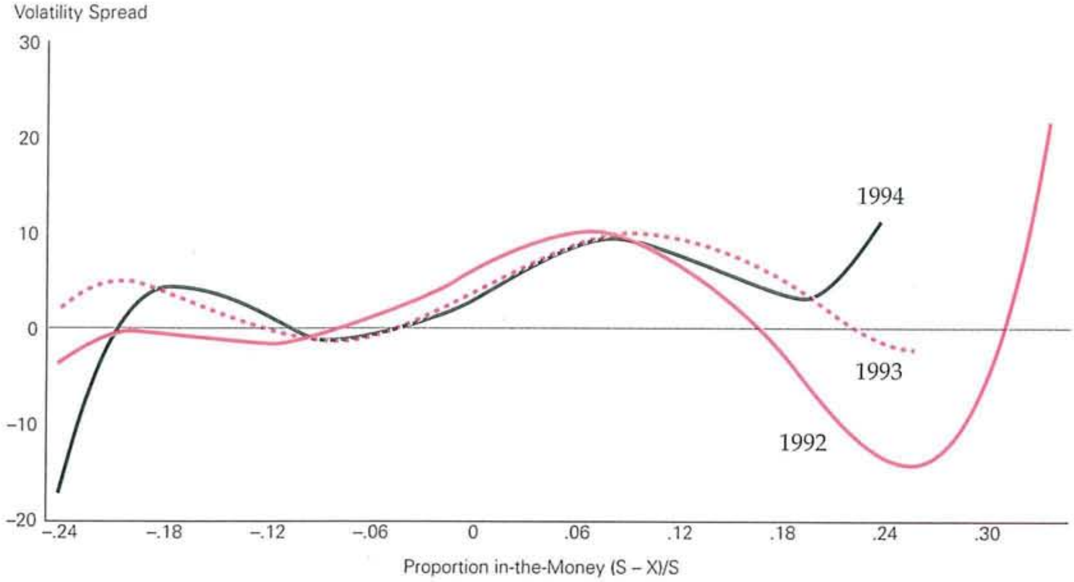

Source: Author's calculations.

Figure 4 shows, for call options only, the pattern of the relationship between the Volatility Spread and the amount by which an option is in-the-money. The vertical axis, labeled Volatility Spread, is the deviation of the implied volatility predicted by the polynomial regression from the group mean of implied volatilities for all options trading on the same day with the same expiration date. For each year the pattern is shown throughout that year's range of values for ITM. While the pattern for each year looks more like Charlie Brown's smile than the standard smile, it is clear that there is a smile in the implied volatilities: Options that are further in or out of the money appear to carry higher volatilities than slightly out-of-the-money options. The pattern for extreme values of ITM is more mixed.

Test 3: A Put-Call Parity Test

Another prediction of the Black-Scholes model is that put options and call options identical in all other respects should have the same implied volatilities and should trade at the same premium. This is a consequence of the arbitrage that enforces put-call parity. Recall that put-call parity implies $P_t + e^{-q(T-t)}S_t = C_t + Xe^{-r(T-t)}$. A put and a call, having identical strike prices and terms, should have equal premiums if they are just at-the-money in a present value sense. If, as this paper does, we interpret at-the-money in current dollars rather than present value (that is, as $S = X$ rather than $S = Xe^{-(r-q)(T-t)}$), at-the-money puts should have a premium slightly below calls. Because an option's premium is a direct function of its volatility, the requirement that put premiums be no greater than call premiums for equivalent at-the-money options implies that implied volatilities for puts be no greater than for calls.

For each trading day in the 1992–94 period, the difference between implied volatilities for at-the-money puts and calls having the same expiration dates was computed, using the ± 2.5 percent criterion used above.[12] Figure 5 shows this difference. While

[12] The expiration dates for the put and call are the same for any day, but on different days the time to expiration of the options will be different.

Figure 5

Difference between Implied Volatilities on SPX Puts and Calls

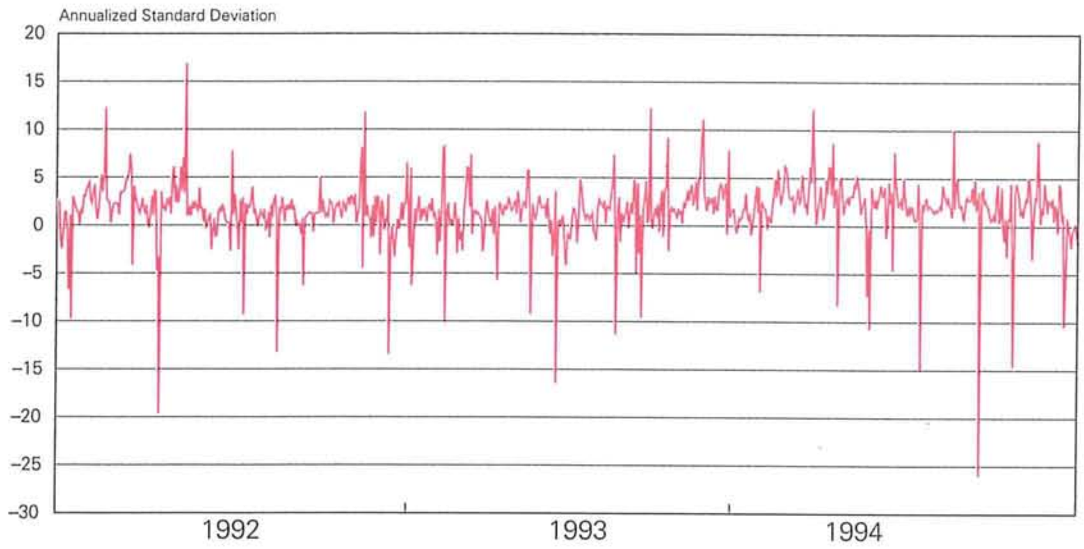

Source: Author's calculations.

puts sometimes have implied volatility less than calls, the norm is for higher implied volatilities for puts. Thus, puts tend to trade "richer" than equivalent calls, and the Black-Scholes model does not pass this put-call parity test.

Test 4: Option Pricing Errors

The tests used thus far have relied on implied volatilities. We now turn to a test based directly on pricing errors, converting information on the range of implied volatilities for repeated trades of SPX options, all traded on the same day and having the same expiration date, into measures of the Black-Scholes pricing error. This test, due in part to Rubinstein (1994), is based on the Black-Scholes model's prediction that repeated trades in the same option over a short interval should, apart from random variations due to chance, reflect the same implied volatility.

The implied volatility data can be converted to information on pricing errors in the following fashion. Suppose we have calculated the implied volatility for each SPX index option traded on a given day and grouped all of those options by their expiration dates. For example, on November 3 there might have been 75 transactions in SPX options that expire on March 17. From this group, select the options with the highest and lowest implied volatility. Let A be the option with the lowest implied volatility and Z be the option with the highest implied volatility, and let a and z denote the respective values of implied volatility.

Now consider Figure 6, which shows the Black-Scholes relationship between pricing errors and the true volatility for these two options. The line AA shows, for each possible level of volatility, the proportional difference between the Black-Scholes value of option A (denoted by C^*) and the observed premium on option A (denoted by C). At volatility level a this difference is zero because the Black-Scholes model predicts $C^* = C$ at option A's implied volatility. Recalling that the Black-Scholes model predicts an increase in the premium as volatility rises, the AA line is upward-sloping. The slope of the line at each point is the option's Vega, which is assumed to be constant in Figure 6. Similarly, the line labeled ZZ is drawn for option Z, the option with the highest implied volatility. The pricing error for option Z is zero at volatility level z, and it increases as volatility increases.

If the Black-Scholes model were always correct,

Figure 6

Option Pricing Error Tests

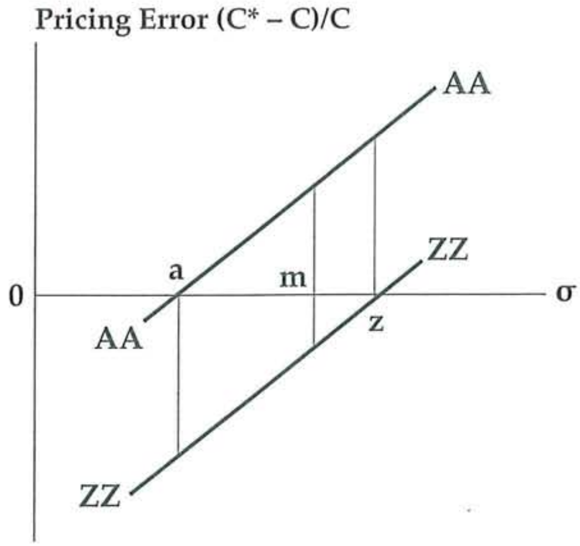

the lines AA and ZZ would coincide because each would have the same implied volatility. The horizontal distance between AA and ZZ is a reflection of the size of the errors in the Black-Scholes model. Of course, we cannot know which option is incorrectly priced. Perhaps both are! But we can calculate the value of the pricing error for each option under different assumptions about the true volatility of the return on the underlying stock.

Assume that option A is correctly priced, so that the true volatility for all options in the group is a. In this case, the pricing error for option Z is shown by the vertical distance from point a to the ZZ line. If, at the other extreme, option Z is correctly priced (the true volatility is z), then the pricing error is the vertical distance to AA at volatility level z. Because we are concerned with the size of the pricing errors, not the sign, we can compute the absolute value of each error just measured. Noting that all options in the group lie on lines between AA and ZZ, so they must have smaller pricing errors at volatility a and at volatility z, we can see that the *larger* of the two absolute errors just measured must set an upper bound on the possible pricing errors. That is, conditional on the assumption that the true volatility is between the lowest implied volatility and the highest implied volatility in a group, the proportional pricing error for all options in the group can be no greater than the largest of the errors measured at the extreme implied volatilities. We call this the Upper Bound error.[13]

We would, of course, like to know what the pricing error is at the true volatility. A rough measure can be obtained by assuming that, for each group of options, the true volatility is the group's average implied volatility. For example, for the hypothetical situation in Figure 6, we can calculate the errors for both option A and option Z at volatility level m, which is the mean implied volatility. The absolute value of the *lowest* of these two errors is called the Lower Bound error. Thus, conditional on the mean implied volatility representing the true volatility, the Black-Scholes model gives an error at least as high as the Lower Bound error.

For any single group of options, the Upper Bound and Lower Bound errors might be a poor measure of the Black-Scholes model's fit. However, if we take the averages over a large number of groups, we can expect a better measure of fit. Table 2 summarizes the results for both calls and puts in each year in the sample and for all three years combined. The combined sample contains 2,034 groups of calls and 2,496 groups of puts. The average number of transactions in each group was 88 calls and 117 puts. The mean time to expiration was 46 days for calls and 58 days for puts.

For the combined sample, the Lower Bound error is 10 percent for calls and 15 percent for puts. This means that if the average implied volatility accurately measures the true volatility, the Black-Scholes model is off the mark by *at least* 10 to 15 percent. This value of the Lower Bound error appears to be stable over time, as shown in the entries for each year. The Upper Bound error is much higher and considerably less stable. For the combined data, the Upper Bound for calls is 97 percent of the observed premium. It is a more moderate 40 percent for puts. Thus, the combined results suggest a 10 to 100 percent error for calls and a 15 to 40 percent error for puts. The conclusion that puts are more accurately priced by the Black-Scholes model is a bit surprising, because the model was originally developed for call options.

[13] There is, of course, some probability that the range of implied volatilities in a group of options does not contain the true volatility. In that event, the Upper Bound error will understate the true error. Unfortunately, we do not know the probability of understatement because we do not know the probability distribution of implied volatilities.

Table 2
Pricing Error Tests

Year	Option Type	Number of Groups[a]	Mean Trades per Group	Mean Term (Days)	Mean Values[b] Lower Bound[b] (Percent)	Upper Bound[b] (Percent)
1992	Call	627	68	45.8	7.7	64.8
	Put	741	86	54.3	13.1	33.8
1993	Call	681	86	45.0	9.7	106.1
	Put	849	104	56.8	17.0	44.0
1994	Call	726	113	47.2	11.4	115.1
	Put	906	154	61.0	15.6	42.1
All 3 Years	Call	2,034	88	46.0	9.7	96.6
	Put	2,496	117	57.6	15.3	40.3

[a]Each group consists of all options of the stated type traded on the same day and having the same expiration date.
[b]The Lower Bound assumes that the true volatility is at the mean of the group's implied volatility. The Upper Bound is the largest error observed in a group.

The importance of these pricing errors depends upon one's perspective. A 10 to 15 percent prediction error is not uncommon for economic data, and the academic economist might feel that a simple abstract model like the Black-Scholes model does quite well with these margins of error. However, for traders and financial practitioners, an error of at least 10 to 15 percent is large enough to drive a truck through. Clearly, these errors are also sufficiently great to drive large research budgets devoted to finding models better than the Black-Scholes model.

Test 5: The Distribution of Stock Prices

Thus far our analysis has focused on the predictions of the Black-Scholes model. Now we look at one of the assumptions of the model: that the instantaneous rate of change in the return on an option's underlying security is log-normally distributed. Formally, the Black-Scholes model assumes that the change in the stock's price is given by a diffusion process in which, at any instant, the rate of change in the stock's price is determined as described in Box 2. This implies that the change in the logarithm of price over a time interval of length T is a normally distributed random variable with an expected value equal to the drift and a variance equal to $T\sigma^2$, where σ is the instantaneous volatility of the rate of return on the stock.

The assumption that stock prices are log-normally distributed has a long history. Made largely as a matter of convenience, it has long been known to be an approximation, sometimes a poor one. Indeed, studies of the frequency distribution of individual stock prices beginning as early as Fama (1965) found that the distribution of changes in the logarithm of stock prices has "fat tails," that is, the relative frequency of very large changes is greater than for the normal distribution. Furthermore, the observed distribution, while having small skewness, is leptokurtic, meaning that it is more bunched in the middle than the normal distribution. The "fat tails" phenomenon has become a stylized fact in finance, and has been used to explain the smile in implied volatilities: Options that are far out of or far in the money have premiums greater than the Black-Scholes prediction. This will show up as higher implied volatilities for off-the-money options than for at-the-money options.

Table 3 reports the descriptive statistics for the daily change in the logarithm of the S&P 500 index for the period January 2, 1980 to March 31, 1995. The statistics are reported for the entire period and for two subperiods: The pre-1987 crash period (January 2, 1980 to September 30, 1987) and the post-crash period (January 4, 1988 to March 31, 1995). Fama and French (1988), among others, have shown that the distribution of daily stock returns is different over weekends than on contiguous trading days. Table 3 reports the descriptive statistics for contiguous trading days and for two-day trading breaks (typically, weekends).

For the entire period, the sample mean and standard deviation for contiguous days (percentages at annual rates) are 21.5 percent and 12.6 percent, while for weekends they are −4.2 percent and 9.1 percent. Thus, in the 1980s and 1990s, the S&P 500 has done well during trading weeks but poorly on weekends. However, this pattern shifted over time. Prior to the '87 Crash, stocks did extremely well during the week and declined over weekends. After the Crash, stocks rose both during the weeks and on weekends, but the performance was particularly good over weekends.

No significant difference is seen between the pre-Crash and post-Crash daily volatility of the S&P 500, either during the week or over weekends. However, the higher moments of the distribution changed

Table 3
Descriptive Statistics for Changes in Log S&P 500[a]

	1/2/80 to 3/31/95	1/2/80 to 9/30/87 (Pre-1987 Crash)	1/4/88 to 3/31/95 (Post-1987 Crash)
Number of days			
Calendar	5,549	2,846	2,639
Trading days	3,859	1,962	1,834
Contiguous trading days	3,030	1,538	1,442
Mean trading days per year			
Total	253	253	253
Contiguous	199	198	199
Return characteristics			
Contiguous trading days			
Mean return (%)	21.48	32.65	8.29
Standard deviation (%)	12.58	12.40	11.27
Skewness	.07	.21	−.94
Kurtosis	7.33	1.80	8.32
Two-day breaks (for example, Friday–Monday)			
Mean return (%)	−4.20	−11.22	20.47
Standard deviation (%)	9.10	6.87	5.73
Skewness	−7.77	−.38	.25
Kurtosis	128.38	.98	2.62

[a]The daily changes in log S&P 500 are (approximately) the daily rate of change in the S&P 500. These were converted into percent at annual rate, as follows: The mean return was calculated on a 365-day calendar year. The standard deviation was calculated on a 199-day trading year for contiguous days, and on 95 days per year for two-day breaks.

dramatically. The skewness of the intra-week daily returns went from positive to negative after the Crash, indicating a shift toward more down days after the Crash. However, the weekend returns turned from negative skewness to positive skewness. In short, after the Crash intra-week returns shifted toward fewer good days, but inter-week returns shifted toward more good days.

In addition to this change in patterns for skewness, the kurtosis of the distribution uniformly increased after the Crash. The kurtosis of stock returns, and the increase in kurtosis, are clear in Figure 7, which shows the relative frequency distributions for both periods as well as the standard normal distribution. Both periods contain many more values in a range within one standard deviation than the normal distribution, as well as slightly more values in the three to four standard deviation range distribution. Thus, the S&P 500 appears to be bunched in the middle range of outcomes, with signs of a few large changes, conforming more to a "thin middle" than to a "fat tails" description.

In summary, the probability distribution of the change in the logarithm of the S&P 500 does not conform strictly to the normality assumption. Not only is the distribution thicker in the middle than the normal distribution, but it also shows more large changes (either up or down) than the normal distribution. Furthermore, the distribution seems to have shifted over time. After the Crash an increase in the kurtosis and a shift in skewness occurred.

An Interim Summary

The analysis of data on almost 500,000 transactions in the SPX call and put options in the 758 trading days of 1992 to 1994 shows abundant evidence against the Black-Scholes model. We find that implied volatility is a poor forecast of actual future volatility, raising doubts about the Black-Scholes assumption that traders are excellent statisticians able to develop optimal forecasts of volatility which are then reflected in option premiums. We find that implied volatilities exhibit a "smile," in contrast to the Black-Scholes model's prediction that implied volatility will be the same for all options having the same underlying stock and the same time to expiration. We find that implied volatilities for at-the-money puts are, other things equal, greater than implied volatilities for at-the-money calls, a result not consistent with the Black-Scholes prediction that put-call parity will ensure that at-the-money puts and calls identical in all respects will have the same premiums and the same implied volatility.

In addition, we have computed estimates of the magnitude of pricing errors. Assuming that the group average of implied volatility is a useful estimate of true volatility, we find that the Black-Scholes model works better for puts than for calls, but that in both cases the errors are economically significant, on the order of *at least* 10 to 15 percent of the actual premium.

Finally, we show that the relative frequency distribution of daily changes in the logarithm of the S&P 500 does not fit the assumption of a normal distribution: It is slightly skewed and highly leptokurtic, and it has signs of fat tails. Thus, the actual distribution of the S&P 500 has slightly more large changes, many more small changes, and fewer mod-

Figure 7

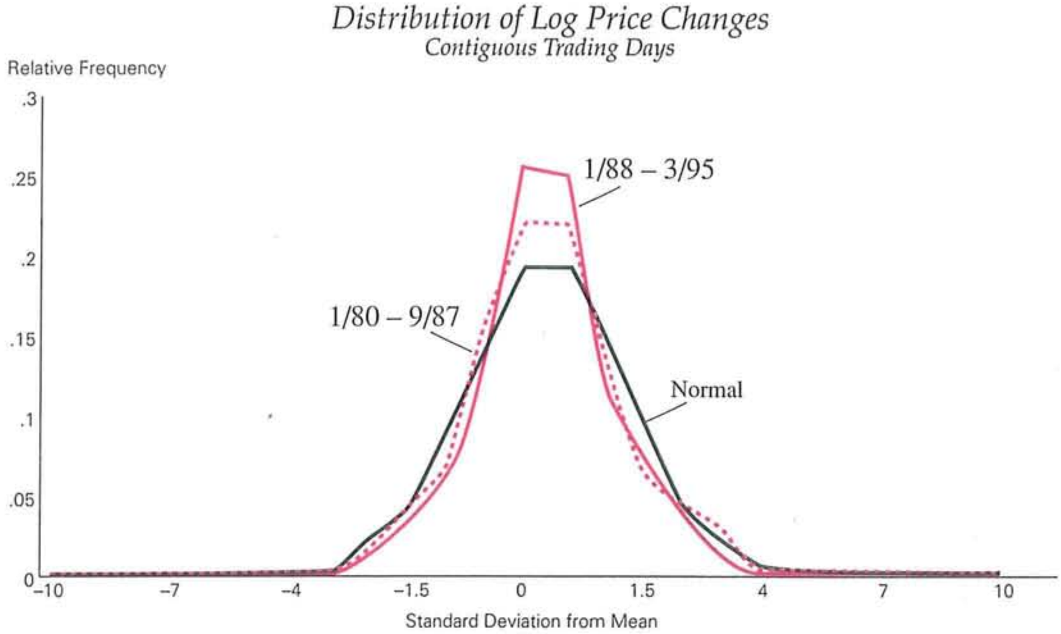

Source: Author's calculations.

erate changes than the normal distribution would predict.

The next section presents some explanations of the failures of the Black-Scholes model.

III. Some Explanations for the Black-Scholes Model Performance

The previous section reports results indicating that the Black-Scholes model provides, at best, a crude approximation to the option premiums observed in data on actual transactions. The reasons for these discrepancies have been the subject of considerable controversy among financial economists, as well as practitioners. This section will examine several reasons why this shortfall might exist.

Limitations on Arbitrage

Underlying the Black-Scholes model of option pricing is an assumption central to finance theory: that traders quickly and efficiently eliminate any discrepancies between actual and theoretical prices by engaging in arbitrage. Two types of arbitrage must be unrestricted to make this happen. First, call options must be fully arbitraged with their replicating portfolios. If a call is underpriced, traders must be able to buy calls and sell the underlying stocks short. If a call is overpriced, traders must be able to buy the stock and write covered calls. Second, arbitrage must enforce put-call parity. If puts are overpriced relative to calls, traders must be able to write puts, buy calls, *and* sell the stock short. If puts are underpriced relative to calls, traders must be able to buy both puts and the underlying stock, and write covered calls.[14]

The put-call parity analysis in the previous section indicates that implied volatilities for puts typically are greater than implied volatilities for calls with identical characteristics, although occasionally call volatilities exceed put volatilities. Stated differently, put options appear to be systematically overpriced relative to identical calls, although occasionally under-

[14] A call is covered if the writer also owns the shares, in which case he can immediately deliver them if the call is exercised. A call is naked if he does not own the shares and might have to buy them under adverse conditions for delivery if the call is exercised.

Box 4: Short Selling and Option Arbitrage

Among the anomalies observed in this paper is that SPX put options appear to be overpriced relative to SPX calls. This is not consistent with the arbitrage conditions underlying option pricing models because, upon observing overpriced puts, traders should simultaneously write puts, sell the underlying stock short, and buy calls. For a number of reasons, arbitrage might not occur in a sufficient volume to correct put overpricing. Note that costs are also associated with correcting underpriced puts, but the barriers to the transactions required (buying both a put and the stock, and writing a covered call) are lower.

- **Restrictions on Entry** In order to sell short, an investor must meet financial standards established by regulation and by brokerage firms. These standards are even higher for investors writing naked options, as an arbitrageur must do in order to take advantage of put overpricing. Thus, some investors will be excluded from the opportunity to engage in the arbitrage required to correct overpriced puts.
- **The Uptick Rule** A short sale can occur only when the price is above the last different price, that is, only on an "uptick." As a result, it might not be possible to simultaneously write a put and sell the stock short. This exposes the arbitrageur to the risk that the price at the short sale's execution is so low that the arbitrage creates a loss.
- **Risks of Premature Termination** The arbitrage required to take advantage of overpriced puts assumes that the position can be held until the expiration of the options. However, a short position can be forced to terminate early if the lender of the shares wants them and replacements cannot be found.
- **Maintenance Margins** Under the Federal Reserve System's Regulation T, the initial margin required for a short sale is equal to the initial margin on a long position; at present, this is 50 percent of the value of the security. The New York Stock Exchange requires member firms to establish a maintenance margin of at least 25 percent, but brokerage houses typically set a maintenance margin of 30 percent or more. Some brokerage houses require higher maintenance margins for short positions than for long positions, reflecting their greater risk exposure for a short sale. This raises the probability of a margin call for a short position.
- **Interest and Dividends** When a stock is purchased, the new owner obtains a right to the dividends and pays the interest rate, either as an opportunity cost or, in the case of a margin purchase, as interest on a security loan. The net cost is $(r - q)$. When a stock is sold short, the short seller is obligated to pay any dividends to the stock lender, but earns no interest on the proceeds of the sale. The net cost is q rather than $(q - r)$. Thus, the sacrifice of interest means that the carrying cost of a short sale is not simply the reverse of the carrying cost of the purchase.
- **Fees for Lending Stocks** Typically, the lending firm loans the shares "flat," that is, without paying or receiving any fees. However, in periods of strong short-selling pressure, the borrowing firm might pay a fee. The borrowing firm can recover this cost through the interest earned on the proceeds of the short sale or by embedding it in the commissions on the short sale.
- **Interest on Unrealized Gains and Losses** Lending brokers typically require a cash deposit equal to 100 percent of the value of the shares. Initially, this is provided by the proceeds of the short sale. If the stock price moves up, the lending broker can ask for additional cash to maintain the 100 percent deposit. This adds to the short seller's margin loan and interest is charged on the debit balance. If the stock price falls, the lending broker releases an amount of cash equal to the unrealized profit, and this is available to earn interest for the short seller. Thus, the short seller pays interest on unrealized losses and receives interest on unrealized gains. The same practice applies to stock purchases on margin.

pricing occurs. This is consistent with limits on arbitrage that prevent correction of overpriced puts but allow correction of underpriced puts. These limits are in two forms, those that inhibit short sales relative to stock purchases, and those that inhibit writing naked options as opposed to buying options or writing call options.

Box 4 summarizes some of the factors that make it particularly risky or expensive to engage in arbitrage to take advantage of overpriced puts. Perhaps the most prominent reason is an asymmetry in costs. When a customer sells stock short, as he must when puts are overpriced, he does not receive interest on the proceeds of the sale. However, if he purchased the shares, he would pay interest in the form of an opportunity cost or in the form of interest on a security loan. This means that a short position is not simply the reverse of a long position, and that short positions carry a higher cost.

Another factor restricting arbitrage involving short sales is the uptick rule, which allows a short sale to be executed only in a rising market. This means that

The relative overpricing of puts might be the result of inhibitions on the arbitrage required to correct put overpricing.

the timing of a short sale might not be synchronized with the option transactions that must also be part of the arbitrage, exposing the arbitrageur to a risk not found in arbitrage involving purchases of shares.

The influence of these short-selling restrictions on put-call arbitrage is exacerbated by limits on writing naked options as opposed to writing covered options. Because the writer of a covered option exposes the brokerage firm to no risk, there are no margin requirements. Naked options, on the other hand, do expose firms to the risk that the writer will not be able to perform if the option is exercised, and brokers require margin protection.[15]

The arbitrage costs just discussed are asymmetrical in that they affect put-overpriced arbitrage more than put-underpriced arbitrage. Other symmetrical arbitrage costs affect both sides equally. Commissions, fees, and the bid-asked spreads for both stock and option trades can be particularly high for arbitraging S&P 500 options contracts if changes in position, either long or short, require transactions in 500 common stocks, some having higher transactions costs than the most actively traded.[16] This consideration limits both sides of the put-call parity arbitrage: Short positions in the S&P 500, required to correct overpriced puts, are expensive, but so are long positions required to correct underpriced puts. The effect of these transactions costs is to limit all arbitrage, not just arbitrage involving short sales.

The existence of transactions costs creates a range within which put-call premiums can vary without eliciting corrective arbitrage. Because the costs of correcting put overpricing exceed those of correcting put underpricing, we should expect some tendency for put overpricing to continue unless it is excessive. Thus, asymmetrical transactions costs can help to explain the relationship we have found in our data: At-the-money puts tend to sell at premiums relative to calls that are greater than allowed by put-call parity.

If the arbitrage-inhibiting factors discussed in Box 4 fully account for the observed tendency toward put overpricing, the failure of put-call parity cannot be laid at the foot of Black-Scholes, for that model abstracts from transactions costs. However, it seems likely that some portion of the observed put overpricing is a true anomaly, in the sense that it reflects an inefficiency in the market for stock index options. The reason for this speculative conclusion lies in the ability of many traders, particularly financial institutions, to arrange their transactions in ways that make the costs of options arbitrage small. For example, short sales of stock indexes can be replicated in the futures markets without the expenses related to selling short through brokers. Furthermore, wholesale traders negotiate the terms with brokerage firms and are not bound by the published terms for retail transactions.

The Non-Normality of Stock Price Changes

As noted above, an important assumption of Black-Scholes—that the price of the underlying security is log-normally distributed—is not validated by the observed distribution of changes in the logarithm of stock prices: The relative frequency distribution is roughly symmetric (with signs of skewness), and very leptokurtic. Furthermore, while there might be evidence of fat tails for individual equities, we observe only minor fat tails for the S&P 500.

One explanation for non-normality, first pro-

[15] Self-regulating organizations set minimum requirements on naked options, subject to Securities and Exchange Commission approval. Brokers are free to set margins higher than the required minimums. One major discount broker requires margin on writing naked equity options equal to the premium received plus 25 percent of the underlying security's market values minus any out-of-the-money amount.

[16] Of course, a trader can choose to arbitrage with a subset of the S&P 500, but this substitutes basis risk for out-of-pocket expenses. In addition, stock index futures contracts are a low-cost way of taking a position on the index.

Box 5: A Jump-Diffusion Model of Stock Prices

Press (1967) first presented a model of stock prices consistent with the non-normality observed by Fama (1965). Press's Compound Events Model was developed further by Merton (1976a, 1976b) and Cox and Ross (1976). According to this model, stock price changes conform to the following "jump diffusion" process.

$$dS/S = \mu dt + \sigma dz + dq \qquad (B5.1)$$

where dz is a Wiener variable (that is, $dz = \epsilon\sqrt{dt}$, with ϵ an independent standard normal random variable), and dq is a random variable defined as $(Y - 1)dt$ if a jump occurs in interval dt and zero otherwise. The variable Y is a jump multiplier for S ($Y = 2$ means a doubling of S), so the proportional change is S is $(Y - 1)$. The jump multiplier, Y, is a random variable assumed to be log-normally distributed, so lnY is normal with mean θ and variance δ^2.

The process controlling whether a jump occurs is assumed to be Poisson, with $\lambda\Delta t$ being the probability of one jump in the small interval Δt and $p(n) = e^{-\lambda h}(\lambda h)^n/n!$ being the probability of n jumps in the interval of length h.

Merton has shown that this jump process results in the following stock-price dynamics:

$$ln[S(t + h)/S(t)] = [\mu - ½\sigma^2 - \lambda\theta]h + [\sigma\sqrt{h}]Z + \Sigma_n lnY_n \qquad n = 0, 1, 2, \ldots \qquad (B5.2)$$

with Z being a standard normal random variable (mean zero and standard deviation 1) and each of the n values of lnY being normally distributed with mean θ and variance δ^2. Thus, the log of the price ratio consists of a drift of $[\mu - ½\sigma^2 - \lambda\theta]h$ plus a normally distributed part consisting of a Wiener variable plus the sum of n "shocks," each being the logarithm of the respective jump size.

Under this jump-diffusion model, the log of stock price-relatives in (B5.2) is a Poisson mixture of normal distributions. Conditional on exactly n jumps occurring, the log of the price-relative will be normally distributed, with mean $[\mu - ½\sigma^2 - \lambda\theta]h$ and variance $[h\sigma^2 + \lambda n(\theta^2 + \delta^2)]$. Because n is a random variable, this variance will be random with a Poisson distribution. Because the expected value of n over an interval of length h is λh, the unconditional variance is $[\sigma^2 + \lambda(\theta^2 + \delta^2)]h$.

Characteristics of the distribution of $ln[S(T)/S(t)]$ over interval $(T - t)$ are determined by the first four cumulants of the distribution, given below:

Cumulants	*Characteristics*
$k_1 = (\mu - ½\sigma^2 - \lambda\theta)h$	Mean $= k_1$
$k_2 = [\sigma^2 + \lambda(\theta^2 + \delta^2)h$	Variance $= k_2$
$k_3 = \lambda\theta(\theta^2 + 3\delta^2)h$	Skewness $= k_3/k_2^{3/2}$
$k_4 = \lambda(\theta^4 + 6\theta^2\delta^2 + 3\delta^4)h$	Kurtosis $= k_4/k_2^2$

The first and second cumulants are the expected value and variance of $ln[S(t{+}h)/S(t)]$. The signs of the third and fourth cumulants determine the directions of skewness and kurtosis. If $\theta \neq 0$, the distribution is skewed, and the sign of the skewness is the sign of θ: When θ is positive, random jumps will increase the stock price on average and the distribution will be positively skewed, having more increases than decreases. The kurtosis will be positive (indicating leptokurtosis) so long as either $\theta \neq 0$ or $\delta > 0$, that is, so long as discrete shocks influence prices.

Thus, the jump-diffusion model is consistent with the observed characteristics of the frequency distribution of daily changes in the logarithm of the S&P 500: leptokurtic (having a thin middle) with a potential for skewness.

The jump-diffusion model also leads to the following theoretical call premium:

$$C = \Sigma_n[e^{-\lambda h}(\lambda h)^n/n!]C_n \qquad \text{with } n = 0, 1, 2, \ldots \qquad (B5.3)$$

stating that the call premium is a weighted average of the conditional Black-Scholes premiums, each conditioned on the number of jumps. Thus, C_n is the Black-Scholes model when there are n jumps, in which case the volatility is the square root of $[\sigma^2 + n(\theta^2 + \delta^2)]$, and $e^{-\lambda}(\lambda)^n/n!$ is the Poisson probability of n jumps.

posed by Press (1967) even before the Black-Scholes paper, is that stock prices are hit by occasional shocks, called jumps, which cause temporary departures from normality. Press's "Compound Events" explanation is that the changes in the log of stock prices conform to a normal distribution in the absence of shocks; this is consistent with the Wiener diffusion process, which underlies Black-Scholes. However, when jumps occur, the distribution changes: It remains normal, but with a variance that depends upon the number of jumps. Thus, the observed distribution is made up of a mixture of different normal distributions, each with a variance depending on the number of jumps. Box 5 discusses the foundations of a "jump diffusion" model and shows that it is consistent with the stylized facts: It results in a relative frequency distribution that *might* be skewed but *will* be leptokurtic.

As shown by Merton (1976a), the jump-diffusion model leads to a specific closed-form model of option prices in which the premium is a weighted average (using the Poisson distribution) of the Black-Scholes premiums for options, each conditional on the number of jumps. Figure 8 shows the theoretical call option values for a Black-Scholes model and a jump-diffusion model. Because the jump-diffusion model incorporates the greater-than-normal proportion of small stock price changes, the jump-diffusion model places greater value on the call in the region around the current stock price. Because the hypothetical option in Figure 8 is at-the-money, the greater value is placed in the region around the strike price.[17]

Figure 8

Jump-Diffusion and Black-Scholes Hypothetical Call Option

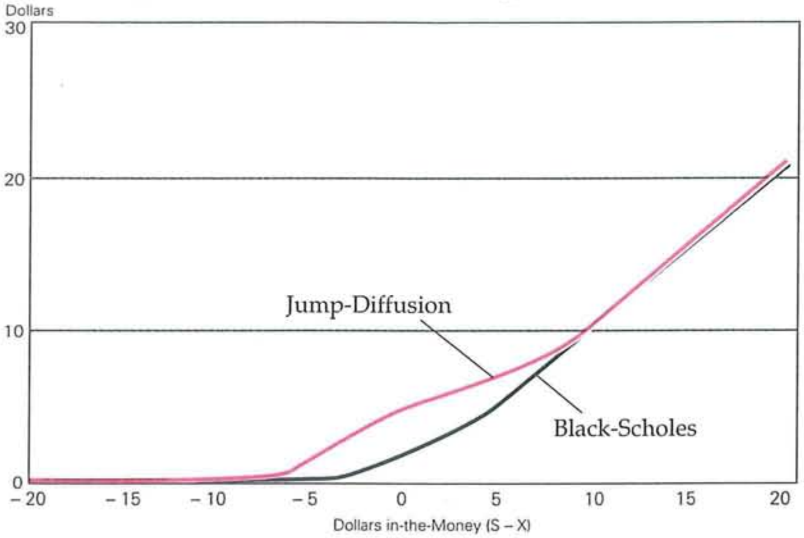

IV. Summary and Conclusions

Recent years have seen a reinvestigation of the "efficient markets" hypothesis (EMH) of financial market performance. An hypothesis that once had widespread acceptance, the EMH has not fared well under newer tests. For example, Fortune (1991) has reviewed the literature on the EMH as it applies to the stock market.

This study suggests that the Black-Scholes model is not consistent with the efficient markets hypothesis. That this is true is also demonstrated by the proliferation of other models of option pricing in recent years, including nonanalytical methods involving numerical analysis. One implication of this finding is that those who are responsible for monitoring financial institutions should not naïvely apply popular formal models of option pricing to assess financial risks. To do so is to invite undoing successful risk management strategies that more informed internal management might adopt.

The goal of this study is to examine the EMH in the context of the market for options on common stocks. The ability of the premier option pricing theory, the Black-Scholes model, to explain observed premiums on S&P 500 stock index options is subjected to a number of tests. The article begins with a summary of the Black-Scholes model, then examines some of the assumptions underlying the model. The second section describes the data used to evaluate the Black-Scholes model and reports the results. Using almost 500,000 transactions on the SPX stock index option traded on the Chicago Board Options Exchange in the years 1992 to 1994, the study finds a number of violations of the Black-Scholes model's predictions or assumptions.

First, the Black-Scholes model assumes that the market forms efficient estimates of the volatility of the

[17] For Figure 8, it is assumed that the jump-diffusion parameters are $\sigma = 0.01$, $\delta = 0.02$ and $\lambda = 0.2$. The Black-Scholes model is the same, but with $\lambda = 0$, that is, no jumps occur.

return on the S&P 500. These estimates then become embedded in the premiums paid for options and can be recovered as the option's "implied volatility." However, several tests of the implied volatility indicate that it is a poor estimate of true volatility: For the SPX contracts, implied volatility is an upwardly biased estimate of the observed volatility. Furthermore, implied volatility does not contain all the relevant information available at the time the option is traded, a violation of the assumption of forecast efficiency.

A second test is based on the Black-Scholes model's prediction that options trades at the same time and alike in all respects except the strike price should exhibit no relationship between the implied volatility and the strike price, or between implied volatility and the amount by which the option is in or out of the money. This study finds, as have other studies, a "smile" in implied volatility: Near-the-money options tend to have lower implied volatilities than moderately out-of-the-money or in-the-money options. A third test is derived from the Black-Scholes model's prediction that put-call parity will ensure that puts and calls identical in all respects (expiration date, strike price, expiration date) have the same implied volatilities. This study finds that puts tend to have a higher implied volatility than equivalent calls, indicating that puts are overpriced relative to calls. The overpricing is not random but is systematic, suggesting that unexploited opportunities for arbitrage profits might exist.

Those who are responsible for monitoring financial institutions should not naïvely apply popular formal models of option pricing to assess financial risks. To do so is to invite undoing successful risk management strategies that more informed internal management might adopt.

A fourth test is based on measures of the pricing errors associated with the Black-Scholes model. Deviations between the theoretical and observed option premiums should be small and random. Instead the study finds systematic and sizable errors. For the entire sample, calls have pricing errors averaging 10 to 100 percent of the observed call premium. Puts appear to be more accurately priced, with errors 15 to 40 percent of the observed premium.

Finally, the distribution of changes in the logarithm of the S&P 500 is examined. The Black-Scholes model assumes that stock prices are log-normally distributed, that is, that the logarithm of the price is normally distributed. It has long been known that this is not true, and the received wisdom is that stock prices exhibit "fat tails" relative to a normal distribution (more extreme changes than the normal distribution would predict). Our analysis of daily values of the S&P 500 confirms a departure from normality for the period 1980 to 1995. We find minor evidence of fat tails, but much evidence of a "too-thin" middle in the distribution—more small changes and fewer moderate-sized changes than the normal distribution would allow.

The paper's third section turns to some explanations for these results. We suggest that the relative overpricing of puts might be the result of inhibitions on the arbitrage required to correct put overpricing. These limitations, in the form of transactions costs and risk exposure, are greater for short selling of stock, the way arbitrageurs would take advantage of put overpricing, than for buying stock, the mechanism for correcting put underpricing.

We also examine a model of stock prices that can explain the non-normality observed in our data. This "jump-diffusion" model is consistent with both the observed skewness and the leptokurtosis in the distribution of stock prices. According to this model, stock prices are usually consistent with a log-normal distribution, but occasional shocks create discrete jumps up or down in the price. This leads to a distribution of stock prices that looks like the one observed for the S&P 500 index, that is, roughly symmetric but with small changes given excessive weight. The jump-diffusion model of stock prices also leads to a specific option pricing model that is a modification of Black-Scholes.

References

Black, Fischer and Myron Scholes. 1973. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy*, vol. 81 (May-June), pp. 637–54.

Canina, Linda and Stephen Figlewski. 1993. "The Informational Content of Implied Volatility." *The Review of Financial Studies*, vol. 6, no. 3, pp. 659–81.

Cox, John and Stephen Ross. 1976. "The Valuation of Options for Alternative Stochastic Processes." *Journal of Financial Economics*, vol. 3 (January/March), pp. 145–66.

Fama, Eugene F. 1965. "The Behavior of Stock-Market Prices." *Journal of Business*, vol. 38, pp. 34–105.

Fama, Eugene F. and Kenneth French. 1988. "Permanent and Temporary Components of Stock Prices." *Journal of Political Economy*, vol. 96 (April), pp. 246–73.

Fortune, Peter. 1991. "Stock Market Efficiency: An Autopsy." *New England Economic Review*, March/April, pp. 17–40.

———. 1995. "Stocks, Bonds, Options, Futures and Portfolio Insurance: A Rose by Any Other Name. . . ." *New England Economic Review*, July/August, pp. 25–46.

Mayhew, Stewart. 1995. "Implied Volatility." *Financial Analysts Journal*, July-August, pp. 8–20.

Merton, Robert. 1973. "Theory of Rational Option Pricing." *The Bell Journal of Economics and Management Science*, vol. 4 (Spring), pp. 141–83.

———. 1976a. "Option Pricing When Underlying Stock Returns are Discontinuous." *Journal of Financial Economics*, vol. 3 (January/March), pp. 125–44.

———. 1976b. "The Impact on Option Pricing of Specification Error in the Underlying Stock Returns." *Journal of Business*, vol. 31, no. 2 (May), pp. 333–50.

Minehan, Cathy E. and Katerina Simons. 1995. "Managing Risk in the '90s: What Should You Be Asking about Derivatives?" *New England Economic Review*, September/October, pp. 3–25.

Press, S. James. 1967. "A Compound Events Model for Security Prices." *Journal of Business*, vol. 40, no. 3 (July), pp. 317–35.

Rubenstein, Mark. 1994. "Implied Binomial Trees." *Journal of Finance*, vol. 64, no. 3 (July), pp. 771–818.

Smith, Clifford. 1976. "Option Pricing: A Review." *Journal of Financial Economics*, vol. 3 (January/March), pp. 3–51.