Option Pricing and the Martingale Restriction

Francis A. Longstaff
University of California at Los Angeles

In the absence of frictions, the value of the underlying asset implied by option prices must equal its actual market value. With frictions, however, this requirement need not hold. Using S&P 100 index options data, I find that the implied cost of the index is significantly higher in the options market than in the stock market, and is directly related to measures of transaction costs and liquidity. I show that the Black-Scholes model has strong bid-ask spread, trading volume, and open interest biases. Option pricing models that relax the martingale restriction perform significantly better.

The no-arbitrage approach to valuing derivative securities has become a standard paradigm in finance. This approach was first introduced by Black and

I am grateful for the comments and suggestions of Kerry Back, Warren Bailey, Giovanni Barone-Adesi, Fischer Black, Michael Brennan, Menachem Brenner, Wolfgang Buehler, Stephen Buser, Linda Canina, Peter Carr, George Constantinides, Ian Cooper, Nicholas Crew, Stephen Figlewski, Margaret Forster, Gunter Franke, Ian Garrett, Thomas George, Mark Grinblatt, Andreas Grunbichler, Bruce Grundy, Campbell Harvey, Michael Hemler, Ayman Hindy, David Hirshleifer, Robert Jarrow, Ravi Jagannathan, Narasimhan Jegadeesh, Andrew Karolyi, Narayana Kocherlakota, Craig Lewis, Robert McDonald, Ehud Ronn, Mark Rubinstein, Barry Schachter, Jose Scheinkman, Eduardo Schwartz, David Shimko, Marti Subrahmanyam, Bruce Tuckman, and workshop participants at Brigham Young University, Cornell University, Dartmouth College, the Kansallis Foundation for Financial Research, INSEAD, the London School of Economics, the University of Lausanne, The Ohio State University, the University of Mannheim, the University of Michigan, Princeton University, Queens University, Salomon Brothers, the University of Texas at Austin, and UCLA, and participants at the 1992 Western Finance Association meetings, the 1993 European Finance Association meetings, the 1994 American Finance Association meetings, the 1994 AMEX Options and Derivatives Colloquium, the 1995 Econometric Society meetings, and for the research assistance of John Wayman and Ravi Jain. I am particularly grateful for the comments of the Editor Andrew Lo and an anonymous referee. All errors are my responsibility. Address correspondence to Francis A. Longstaff, Anderson Graduate School of Management, UCLA, 405 Hilgard Avenue, Los Angeles, California 90095.

The Review of Financial Studies Winter 1995 Vol. 8, No. 4, pp. 1091–1124

Scholes (1973) and has been extended by Cox and Ross (1976), Harrison and Kreps (1979), Harrison and Pliska (1981), Merton (1973), and others. In its most general form, the no-arbitrage approach is often called risk-neutral valuation.

In the no-arbitrage framework, the price of an option is given by taking the expectation of its payoff with respect to a risk-neutral or certainty-equivalent density, and then discounting the expectation at the riskless rate. To avoid arbitrage, however, the mean of the risk-neutral density must satisfy the martingale restriction. This restriction is simply that the price of the underlying asset implied by the option pricing model must equal its actual market value. Harrison and Kreps (1979) show that in frictionless markets, the violation of this fundamental restriction implies the existence of arbitrage opportunities.

When there are transaction costs or other market frictions, however, the martingale restriction need not be satisfied. This is because the no-arbitrage framework is only able to place bounds on option prices when there are market frictions. In this case, option prices are determined by equilibrium rather than no-arbitrage considerations. If market frictions are significant, then imposing the no-arbitrage martingale restriction on a model such as the Black-Scholes may limit its ability to explain option prices.

In this article, I examine whether the martingale restriction holds for an extensive sample of actively traded S&P 100 index option prices. This analysis is important for several reasons. In particular, the results provide direct evidence about the effects of transaction costs, market illiquidity, or other types of frictions on option prices. Furthermore, examining whether the martingale restriction holds provides a simple new approach for testing no-arbitrage option pricing models. Finally, this analysis can identify key factors omitted by no-arbitrage models which should be incorporated into more general models of derivative security prices.

Using daily option data, I invert the Black-Scholes model and solve simultaneously for the implied index value and volatility. I then compare the implied index value directly to the actual index value. The martingale restriction is strongly rejected by the data. I find that the implied index value exceeds the actual index value for more than 99 percent of the sample. On average, the implied index value is roughly one-half percent higher than the actual index value. Because an option can be viewed as a levered position in the underlying asset, these results suggest that it is more expensive to purchase stock via the options market than in the stock market.

To determine whether the rejection of the martingale restriction is in fact due to the presence of market frictions, I regress the differences between the implied and actual index values on a number of transaction cost and market liquidity variables. I find that an increase in the average option bid-ask spread increases the implied cost of the index in the options market. Similarly, the implied cost of the index decreases when the options market becomes more liquid. These results indicate that the Black-Scholes model has strong bid-ask spread, trading volume, and option interest biases. I show that these are distinct from the previously documented biases of the Black-Scholes model. These results suggest that transaction costs and liquidity effects play a major role in the valuation of index options.

I also examine how the pricing performance of the standard Black-Scholes model compares with that of an 'equilibrium' version of the model in which the martingale restriction is relaxed. I show that more than half of the pricing error of the Black-Scholes model is eliminated by allowing the cost of the index in the options market to differ from the actual market value. In addition, relaxing the martingale restriction eliminates most of the biases of the Black-Scholes model. I also find that the standard Black-Scholes model results in upward-biased estimates of implied volatility. These results indicate that option pricing models which incorporate the effects of market frictions have the potential to significantly improve upon the performance of no-arbitrage option pricing models.

I also examine whether the martingale restriction is satisfied by a number of other no-arbitrage option pricing models. The results are similar to those for the Black-Scholes model and are consistent with the interpretation that transaction costs and market liquidity are reflected in option prices.

The remainder of this article is organized as follows. Section 1 reviews the risk-neutral valuation framework and the martingale restriction. Section 2 describes the data. Section 3 presents the tests of the martingale restriction. Section 4 examines the properties of the differences between implied and actual index values. Section 5 compares the traditional Black-Scholes model with an unrestricted version of the model. Section 6 considers a number of alternative no-arbitrage option pricing models. Section 7 summarizes the paper and discusses the results.

1. The Martingale Restriction

All no-arbitrage option pricing models impose a common restriction on option prices. This is simply the requirement that the implied cost of the underlying asset in the options market must equal its actual market value. I designate this requirement the martingale restriction. In this section, I review the basic no-arbitrage risk-neutral valuation framework and show why the martingale restriction can be viewed as its primary empirical implication. In the absence of transaction costs, illiquidities, or other market frictions, the martingale restriction must hold exactly in order to avoid arbitrage opportunities. If there are market frictions, however, the martingale restriction need not hold.

The origin of the risk-neutral valuation model is in the Black and Scholes (1973), Cox and Ross (1976), and Merton (1973) no-arbitrage theory of option pricing. Harrison and Kreps (1979) model this valuation framework more formally and introduce the notion of a pricing functional which operates on the payoff function for a contingent claim. This pricing functional transforms the payoff function into a price which is consistent with the underlying asset price in the sense of avoiding arbitrage opportunities.

To illustrate the role of the martingale restriction, fix two dates $t = 0$ and $t = T$, and consider the valuation at time zero of derivative securities with payoffs at time T. Let (Ω, B, P) denote the underlying probability space defining the possible realizations X_T of values for the underlying asset at time T. Here, Ω represents the set of possible outcomes of X_T, B represents a σ-algebra of sets in Ω, and P denotes the underlying probability measure that assigns probabilities to the various elements of B. I restrict attention to the set of contingent-claim payoffs that are B-measurable and square integrable with respect to P. I denote this space of contingent claims $L^2(\Omega, B, P)$.

Harrison and Kreps (1979) show that the pricing operator mapping time-T payoff functions into time-zero prices must have certain properties in order to avoid the possibility of arbitrage opportunities. Specifically, the pricing operator must be linear, continuous, and strictly positive. Furthermore, the pricing operator, when applied to X_T, must give the current price of the underlying asset X_0. Intuitively, the reason why these properties are necessary is clear. The linearity requirement ensures that the pricing operator has the portfolio property—that the price of a portfolio is the same as the sum of the prices of its components. The continuity requirement implies that small changes in the payoff function result in small changes in the price of the contingent claim. The positivity requirement means that positive payoffs map into positive prices. The condition that X_T maps into X_0 simply requires the pricing functional to be internally consistent. If any of these four conditions are violated, the possibility of generating arbitrage profits from the pricing distortions exists.

Given these basic properties for the pricing functional, the Riesz representation theorem for L^p spaces provides a simple characterization of the pricing functional as an expectation operator. Following Harrison and Kreps, the price of a contingent claim with payoff $F(X_T)$, where $F(X_T) \in L^2(\Omega, B, P)$, is given by

$$E_P[\rho F(X_T)], \tag{1}$$

where $\rho > 0$, $\rho \in L^2(\Omega, B, P)$, and E_P is the expectation operator associated with P.[1] This Riesz representation of the pricing functional is consistent with many asset-pricing models. For example, if ρ is interpreted as the intertemporal marginal rate of substitution for a representative agent, then Equation (1) becomes a standard Euler equation and is compatible with models such as Breeden (1979), Constantinides (1989), Cox, Ingersoll, and Ross (1985a), Hansen and Richard (1987), Hansen and Singleton (1983), Lucas (1978), and Rubinstein (1976).

Let D_T represent the time-zero price of a unit discount bond with maturity date T. From Equation (1),

$$D_T = E_P[\rho]. \tag{2}$$

Assuming D_T is bounded above zero for all T, Equation (1) can be rewritten as

$$D_T E_P[\rho F(X_T)/D_T]. \tag{3}$$

Note that Equation (3) holds even if interest rates are stochastic. Since ρ / D_T is positive, square integrable, and has an expected value of one, the Lebesgue-Radon-Nikodym theorem can be applied to simplify the representation of the pricing operator further:

$$D_T E_Q[F(X_T)], \tag{4}$$

where E_Q is the expectation operator relative to a new probability measure Q defined on the probability space (Ω, B, Q). The measure Q is equivalent to P in the sense that Q assigns probability zero to a set in B if and only if P assigns probability zero to the same set.[2]

From Equation (4), contingent-claim values are given by taking the expectation of the payoff with respect to Q, and then discounting at the riskless rate—as if market participants were risk neutral. For this reason, Equation (4) is termed the risk-neutral valuation model and Q is designated the risk-neutral pricing measure. Note that the risk-neutral valuation model can also be viewed as a certainty-equivalent model.

Observe that the risk-neutral valuation model is different from the risk-adjusted valuation model. In the risk-adjusted valuation model,

---

[1] Expectations obtained by applying E_P are conditional on all prices and state variables that generate the σ-algebra B.

[2] This condition is known as absolute continuity and follows because $\rho / D_T > 0$. We assume that P and Q are absolutely continuous with respect to Lebesgue measure on the real axis. This implies that P and Q have density functions.

contingent-claim prices are given by

$$E_R\left[\exp\left(-\int_0^T r(s)\,ds\right)F(X_T)\right], \tag{5}$$

where R is a risk-adjusted probability measure equivalent to P, and r is the short-term interest rate. Although the risk-neutral and risk-adjusted valuation models give the same prices for contingent claims, the pricing measures Q and R are the same only when interest rates are not stochastic.

To see that the pricing functional in Equation (4) preserves the basic properties of viable pricing operators, recall that the expectation operator is linear and that probability measures are positive. This guarantees that the pricing operator has the portfolio property and that positive payoff functions map into positive prices. Furthermore, the boundedness of the expectation operator ensures that Equation (4) has the continuity property.

It is important to recognize, however, that the linearity, positivity, and continuity properties hold for any choice of Q. Thus, these three properties alone are not sufficient to give empirical content to the valuation model in Equation (4). The only property not guaranteed by the representation of the pricing operator as a certainty equivalent is that the pricing operator gives the current price of the underlying asset when applied to X_T.

To close the model and satisfy the remaining no-arbitrage condition, I require that the probability measure Q have the property

$$X_0 = D_T E_Q[X_T]. \tag{6}$$

This requirement is a simple restriction on the mean of the probability measure Q and is analogous to the first moment restrictions imposed by other asset pricing models such as the CAPM or the APT. Intuitively, this restriction means that the price of the underlying asset implicit in the derivatives market must equal the actual market value of the underlying asset. If markets are frictionless, then the violation of this condition implies the existence of a riskless arbitrage opportunity. I designate this condition the martingale restriction.[3]

In markets with transaction costs or other frictions, however, the no-arbitrage conditons are not sufficient to price options and the martingale restriction need not hold. For example, Levy (1985), Perrakis and Ryan (1984), and Ritchken (1985) show that when there are transaction costs or other frictions, the no-arbitrage conditions only place

[3] When interest rates are stochastic, Q is defined only for a specific horizon T. Hence, the marginale restriction in Equation (6) is only required to hold for a specific horizon T.

bounds on option prices. Boyle and Vorst (1992) and Leland (1985) show that when transaction costs are incorporated into the analysis, the value of a replicating portfolio for an option can be expressed as a discounted expectation or certainty equivalent. The expectation, however, must be taken with respect to a path-dependent probability measure which will generally not satisfy the martingale restriction. When there are transaction costs, market illiquidities, or other frictions, options must be priced by equilibrium rather than no-arbitrage conditions.

A major implication of this is that examining whether the price of the underlying asset implied by option prices equals the market value of the underlying asset can provide information about whether market frictions are reflected in the pricing of options. In addition, this suggests that if violations of the martingale restriction are observed, they should be related to variables proxying for transaction costs, option liquidity, or other types of market frictions. These implications provide the motivation for the empirical tests conducted in this article.

2. The Data

The prices used in this study are for the S&P 100 index options traded at the Chicago Board Options Exchange (CBOE). Since their introduction in 1983, these options have experienced dramatic growth in popularity and are now one of the most actively traded option contracts in the world.

The S&P 100 index options are cash settled and are listed on a monthly expiration date cycle. Options with expiration dates in the three nearby months represent the majority of trading volume. Exercise prices are set at five-point intervals to bracket the current value of the underlying S&P 100 index. Option prices are expressed in terms of dollars and fractions per unit of the S&P 100 index. Each point represents \$100. The minimum fraction is 1/16 for options trading below 3, and 1/8 for all other options.

The data for the study were obtained from the CBOE Market Data Retrieval tape and include all last-sale transactions and bid-ask quotations during 1988 and 1989 for all S&P 100 index options. The bid-ask quotations are reported by CBOE employees who are physically located among the roughly 400 traders on the trading floor. All quotes are for a trade size of 10 contracts. Quotes may be recorded as frequently as 30 times a minute for actively traded options.

In examining the martingale restriction, I use data for call options only. The reason for this is that S&P 100 index options have an American exercise feature that allows the options to be exercised prior to maturity. As shown by Merton (1973), the early exercise of put options can be optimal even if the underlying asset does not pay dividends. Thus, the value of an American put will generally exceed that of a European put. In contrast, Merton shows that if the underlying asset pays a continuous stream of fixed dividends, then early exercise of an American call may not be optimal. Brenner, Courtadon, and Subrahmanyam (1987) demonstrate that the small relative size and roughly continuous nature of the dividend stream on the S&P 100 makes the Merton result applicable to S&P 100 index call options. Thus, the American exercise feature should have little effect on the prices of S&P 100 index call options.[4] By using call option data only, I mitigate the possibility of the American exercise feature biasing the test results.

The call prices used in the sample are drawn from the universe of prices by the following procedure. First, I restrict our attention to the five-minute window from 2:00 P.M. to 2:05 P.M. Using data from the same period each day allows us to avoid the possibility of intraday effects in the S&P 100 index options market affecting the results. I use this time frame to avoid data drawn from periods near the market opening at 8:30 A.M., the low-volume midday period, and the market closing at 3:15 P.M. I then take the midpoint of the first bid-ask price quotation given for each option during the window. I use bid-ask prices rather than transaction prices for several reasons. First, as shown by George and Longstaff (1993) and Phillips and Smith (1980), the bid-ask spread can represent a significant proportion of the value of an option. For example, the bid-ask spread for an out-of-the-money S&P 100 index call option is often as large as 30 percent of the midpoint value of the option. Clearly, inferences about option-pricing models based on transaction data could be affected by whether the transaction was at a bid or an ask price. Secondly, transaction prices could also be affected by the size of the transaction executed. The advantage of using bid-ask prices is that quotes are for a standard-sized trade of 10 contracts. Furthermore, CBOE rules require that quotations made by S&P 100 index option market makers be firm for 10 contracts—the bid-ask prices quoted represent actual prices at which transactions could be executed. Call options that do not have a bid-ask quote during five-minute window are excluded from the sample for that day.

This procedure results in a set of virtually simultaneous option prices for each day in the sample period, where the options vary in terms of their strike prices and expiration dates. I require that there

---

[4] Empirical evidence by Harvey and Whaley (1992) shows that the average value of the early exercise premium in at-the-money American calls on the S&P 100 index is about 2 to 3 cents, representing less than one-quarter percent of the total value of the call option.

be prices for at least four calls available for each day included in the sample. Fewer than 10 days were excluded from the sample period because of this criteria.

Each bid-ask quotation record in the sample includes the value of the S&P 100 index as of that time. Since the index is updated continuously, this ensures that the bid-ask quotation and the index value included in the record are virtually simultaneous. As the estimate of the index value during the five-minute window, I use the average of index values reported in the records for each of the options for that day. In general, however, there is little if any difference in the timing of the option prices and index values used.

Although the dividend stream associated with the S&P 100 index does not lead to significant early exercise premia in the call options, the dividend stream affects the analysis in another way. Intuitively, this is because the underlying asset for the call option is not actually the S&P 100 index, but the S&P 100 index minus the present value of all dividends to be paid prior to the expiration of the option.

To make the dividend adjustment, I obtain the actual dividends on the S&P 100 index for each day during the sample period. The data are obtained from Standard and Poors *100 Information Bulletin.* Using the actual dividends on the S&P 100 index results in more precise tests than using an average dividend rate. I then obtain estimates of the term structure for each day in the sample period and use the appropriate yields to discount each of the dividends received during the remaining life of each option.[5] Thus, our estimates of the present value of dividends reflect the actual amount of the dividends, their timing, and the actual discount factor for the dividend. The values of D_T used in the tests are based on the corresponding maturity Treasury-bill yields reported in the *Wall Street Journal.*

Finally, I eliminate from the sample any set of call prices that violates one of Merton's (1973) distribution-free bounds. In particular, I eliminate sets that violate the upper or lower boundary conditions, the convexity relation, or the restriction on the difference between call prices divided by the difference in their strike prices. These filters ensure that there are no static arbitrage opportunities in the data set. The total number of options excluded for violating one of these static arbitrage bounds is less than one-half percent of the total number of options in the sample. In most of these cases, the arbitrage is on the order of 10 cents, which is likely smaller than the transaction costs associated with implementing the arbitrage strategy.

[5] Because 1990 daily dividend data was not available to us, the last day included in the sample is November 16, 1989.

Table 1
Summary statistics for the S&P 100 index options included in the sample

	Mean	Std. Dev.	Min.	Median	Max.
Number	10.32	3.78	4.00	11.00	20.00
Index value	273.80	27.41	233.18	262.88	334.08
Moneyness	−1.97	3.98	−14.42	−1.99	19.50
Avg. time to exp.	28.27	9.72	2.00	28.00	55.00
Avg. bid-ask spread	0.153	0.036	0.078	0.151	0.331
Call volume	102,553	37,964	18,544	94,396	230,208
Open interest	323,955	64,513	160,983	322,804	545,497

The sample consists of 444 daily sets of call options, where each set consists of prices for all available call options with bid-ask quotations during the five-minute window beginning at 2:00 P.M. Number is the number of call prices included in the sample for a given day. Moneyness is the difference between the index value and the average strike price of the options included in the sample for a given day. Average time to expiration for the options is expressed in days. Call volume is the total reported trading volume of all S&P 100 index call options for a given day. Open interest is the total open interest of all S&P 100 index call options for a given day. The sample period is January 1, 1988, to November 16, 1989.

The resulting data set includes option prices for 444 days during the sample period. The number of call prices available on a given day ranges from 4 to 20, with a median of 11. In addition to the option prices, the data set includes the following information for each daily observation: the corresponding S&P 100 index value, the present value of dividends to be paid during the life of each option, the present value of one dollar to be received at the expiration date of each option, the average bid-ask spread of the options, and the total trading volume and open interest for all S&P 100 index calls for that day. Summary statistics for the data are given in Table 1.

3. Testing the Maritngale Restriction

In this section, I test the martingale restriction imposed by the Black-Scholes model by examining whether the value of the S&P 100 index implied from option prices equals the actual value of the index.

3.1 The empirical results

The empirical approach is a simple one. Using all of the option prices available during the five-minute window for a given day, I invert the Black-Scholes model to estimate both the implied index value and the implied volatility. Since the number of option prices ranges from 4 to 20, it is generally not possible to find a single implied index value and volatility estimate that exactly fit all of the call prices. Consequently, these two parameters are estimated via gridsearch by minimizing the sum of squared deviations between the theoretical and actual option prices. This procedure is repeated for each of the 444 daily sets of option prices in the sample.

Once the implied index value is estimated, it can then be compared directly to the actual index value. I focus on the percentage pricing difference between the implied and actual index values. This percentage difference is defined as the difference between the implied and actual index values divided by the actual index value.

From the earlier discussion, it is clear that estimating the implied index value and volatility is the same as estimating the first and second moments of the risk-neutral density. Since the first and second moments completed specify the lognormal risk-neutral density implied by the Black-Scholes model for horizon T, our approach parallels other research which focuses on inferring the risk-neutral density from option prices. Examples of this include Banz and Miller (1978), Breeden and Litzenberger (1978), Hutchinson, Lo, and Poggio (1994), Rubinstein (1994), and Shimko (1993). In a sense, our approach can also be viewed as a simple extension of the familiar technique of inverting option prices to solve for the implied second moment of the pricing density. Research focusing on implied volatility includes Canina and Figlewski (1993), Chiras and Manaster (1978), Latane and Rendleman (1976), and Schmalensee and Trippi (1978).

Related work includes Manaster and Rendleman (1982) who invert sets of option prices for individual stocks to solve for the implied stock price and volatility parameter. Although they focus more on the issue of whether the implied stock price is useful in predicting returns and include only longer maturity options in their sample, they find some evidence that the implied stock price is higher than the actual stock price. Other related work includes Fackler (1986), Fackler and King (1990), Madan and Milne (1994), Sherrick, Irwin, and Forster (1990, 1992).[6]

The empirical results are reported in Table 2. As shown, the martingale restriction imposed by the Black-Scholes model is strongly rejected by the data. The percentage difference between the implied index value and the actual index value is positive for 442 of the 444 observations. The z-statistic for the hypothesis that positive and negative differences are equally likely is 20.89. The mean percentage difference is 0.465 with a t-statistic of 31.52. The median percentage difference is 0.410.

Table 2 also reports the results by the number of option prices used in estimating the implied index value. As shown, inferences about the martingale restriction are the same for all of the categories. The mean percentage differences range from 0.406 to 0.628 for the various

[6] Grundy (1991) and Lo (1987) examine the relation between the moments of the original density function for the underlying asset and the distribution of option returns. Bates (1991) estimates the probability of a jump in the value of the underlying asset from option prices.

Table 2
Summary statistics for the estimates of the percentage pricing differences between the implied value of the index and the actual index value

Number of calls	Proportion positive	z-statistic	Mean	Standard deviation	t-statistic	Minimum	1st quartile	Median	3rd quartile	Maximum	N
4–5	1.000	8.12	0.509	0.356	11.71	0.017	0.280	0.430	0.604	1.887	66
6–7	1.000	7.14	0.628	0.318	14.12	0.207	0.419	0.530	0.745	1.618	51
8–9	1.000	7.62	0.440	0.295	11.36	0.000	0.199	0.385	0.596	1.117	58
10–11	1.000	9.80	0.433	0.296	14.35	0.015	0.217	0.375	0.561	1.428	96
12–13	0.989	9.11	0.445	0.247	16.80	−0.019	0.269	0.416	0.533	1.426	87
14–20	0.988	9.06	0.406	0.317	11.88	−0.58	0.212	0.341	0.506	1.532	86
All	0.995	20.89	0.465	0.311	31.52	−0.058	0.254	0.410	0.579	1.887	444

Proportion positive is the fraction of observations that result in a positive estimate of the pricing difference. The z-statistic tests the hypothesis that the proportion of positive estimates is 0.50 and is distributed as a standard normal variate. The t-statistic tests the hypothesis that the mean estimate is zero. Number of calls denotes the number of call prices available on a given day from which the daily estimate of the index value is implied. N denotes the number of daily observations with the indicated number of calls. The sample consists of 444 daily observations.

categories, and the t-statistics for the means are all in excess of 11. The uniformity of the results across the various categories strongly suggests that the results are not an artifact of the number of options used in estimating the implied index value.

Recall that a call option can be viewed as a levered position in the underlying asset. Intuitively, these results imply that it is more expensive to purchase stock via the options market than directly in the stock market. There are several possible reasons why this cost may be higher in the options market. For example, the higher cost may simply reflect the higher transaction costs in the options market than in the stock market. This is consistent with George and Longstaff (1993) who find that the average bid-ask spread for a share of stock synthesized by options is roughly twice as large as the average bid-ask spread for NYSE stocks.

Similarly, Boyle and Vorst (1992), Leland (1985), and others show that the transaction costs associated with dynamic trading strategies that synthesize option payoffs can be economically significant. In equilibrium, the present value of these costs may be reflected in the market prices for these options. In addition, the liquidity of the options market may affect the implicit valuation of the equity component of an option. Several recent papers addressing the relation between transaction costs and the equilibrium valuation of securities include Amihud and Mendelson (1986), Constantinides (1986, 1993), Tuckman and Vila (1992), and Vayanos and Vila (1992).

3.2 Diagnostic tests

To ensure that these results are robust, it is important to examine the sensitivity of the estimates of the implied index value to alternative empirical specifications. In this section, I report the results from several alternative specifications in order to provide diagnostic checks on these results.

In Table 2, I use the midpoint of the bid-ask spread as the point estimate of the option price. Since the implied index value is virtually always higher than the actual index value, it is possible that using a lower value for the point estimate of the option price could change the results. Accordingly, I reestimated the implied index values using the bid price for each option.

The results are almost identical to those reported in Table 2. The mean percentage difference is 0.512 with a t-statistic of 34.28. The median percentage difference is 0.457. Of the 444 daily estimates of the implied value of the index, all 444 are positive.[7] Similar results are

[7] Intuitively, it may seem that using the lower bid price should result in a smaller estimate of the

obtained when the percentage differences are estimated using ask prices for the options. Thus, the results are not due to the choice of the midpoint of the bid-ask spread as the point estimate of the call price.

In inverting the Black-Scholes formula, I subtract the present value of the dividends to be paid during the life of the option from the estimated index value. This adjustment for dividends is described in Black (1975), Gibson (1991), and others. In doing this, however, I implicitly assume that the dividends are known with certainty at the date the option is valued by the market. If dividends are not known with certainty, then the empirical results may be biased. In particular, the empirical estimates of the implied index value will be upward biased if the expected dividend is less than the actual dividend.

As a diagnostic check, I reestimate the implied index value under the extreme assumption that the market expects no dividends at all during the life of the option. Even with this extreme assumption, the empirical results are similar to those reported in Table 2. The mean value of the percentage difference is 0.305 with a t-statistic of 24.29. The median percentage difference is 0.250. Of the 444 daily estimates of the implied index, 416 are positive. These results demonstrate that the rejection of the martingale restriction is not due to dividend uncertainty.

Although the early exercise premium in S&P 100 index call prices is small, it may still have an effect of the estimation of the implied index value. In particular, the early exercise premium may lead to a higher estimate for the implied index value since the value of the call option is increasing in the mean of the risk-neutral density. To check this possibility, I subtract from each call option price the estimated amount of the early exercise premium and reestimate the implied index value. The estimates of the early exercise premium are based on the mean values of the early exercise premiums for S&P 100 index calls reported in Harvey and Whaley (1992) during the same sample period as this study. Once again, the empirical estimates of the implied index values are very similar to those in Table 2. The mean value of the percentage difference is 0.448 with a t-statistic of 30.48. The median value is 0.394. Only 4 of the 444 daily estimates are negative. These results indicate that the rejection of the martingale restriction is not due to the American exercise feature of the S&P 100 index call options.

implied index value. What actually happens is that the implied volatility estimate is smaller and the implied index value is slightly higher when the bid price is used. Thus, using the bid price rather than the midpoint effects both the first and second moments instead of just the first or the second.

4. Properties of the Pricing Differences

In addition to comparing the implied index value to the actual index value, it is important to examine the empirical properties of the pricing differences. If the violations of the martingale restriction are due to the effects of market frictions on option prices, then the difference between the implied and actual index values should be related to measures of transaction costs and option market liquidity.

Figure 1 plots the time series of percentage differences between the implied index value and the actual index value. As shown, there is considerable time series variation in the percentage differences. The percentage differences are generally highest at the beginning of 1988. This is the period immediately after the 1987 stock market crash. The percentage differences decline significantly by the second half of 1988 and remain at lower levels through most of 1989. Immediately after the October 13, 1989, minicrash, however, the percentage differences increase to levels similar to those at the beginning of 1988. This time-series variation suggests that violations of the martingale restriction may also be related to market events such as the recent path of stock index prices.

To examine these implications, I regress the percentage pricing differences on variables reflecting the transaction costs of options, option market liquidity, and recent stock market movements. In these regressions, I also control for the other pricing biases of the Black-Scholes model. This ensures that the market friction variables are not simply proxying for previously documented biases of the Black-Scholes model.

A number of previous studies have examined the pricing biases of the Black-Scholes model. Examples of these studies include Chiras and Manaster (1978), Macbeth and Merville (1980), and Rubinstein (1985). These studies generally find evidence of three types of bias: a time to expiration bias, a moneyness bias, and a volatility bias. To control for these biases, I include the average time to expiration and moneyness (index value minus strike price) of the options as independent variables in the regression. To control for volatility, I include the current and first two lagged values of the absolute daily return on the index in the regression. Intuitively, this allows volatility to be represented as a linear combination of recent absolute returns and is similar to an ARCH model. I use this simple proxy for volatility rather than the implied volatility estimate since implied volatility is estimated jointly with the implied index value, which could induce a spurious correlation simply because of sampling variability.

As a measure of the transaction costs of the options, I use the average bid-ask spread for the options used in computing the implied

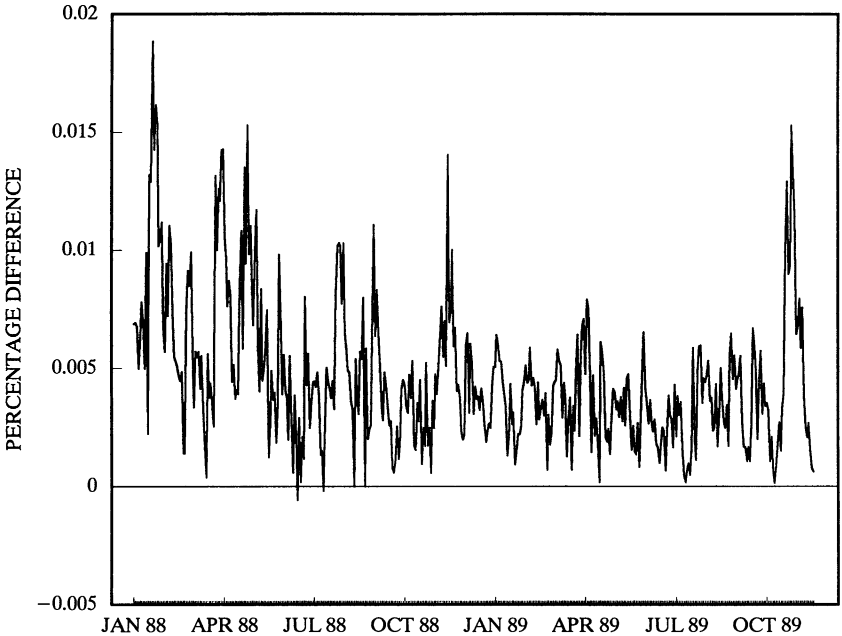

Figure 1
Percentage difference between the implied value of the index and the actual market value of the index
The implied value of the index is obtained by inverting the Black-Scholes model using daily sets of simultaneous S&P 100 index call option prices.

index value. As shown by George and Longstaff (1993), there is considerable cross-sectional variation in the bid-ask spreads for S&P 100 index options. Furthermore, these bid-ask spreads are directly related to the market-making costs and risks faced by market participants.

I use several related measures of market liquidity in the regressions. As a proxy for the total demand for call options, I include the total open interest of all S&P 100 index call options. As one measure of trading activity, I use the total trading volume for all S&P 100 index call options. As another measure of trading activity, I use the total number of calls used to compute the implied index value. Recall that this number reflects the number of calls for which quotes are available during the five-minute window each day. Thus, this number provides a direct measure of market liquidity. As shown in Table 2, there is little or no evidence of a univariate relation between the percentage differences and the number of options. As a proxy for liquidity, however, the number of options could still have explanatory power in a multiple regression specification.

Finally, to capture the possibility of market-related or path-dependent effects on option pricing similar to those suggested by Boyle and Vorst (1992) and Leland (1985), I include current and lagged daily returns on the S&P 100 index as explanatory variables.[8] In estimating the regressions, I use a Cochrane-Orcutt procedure to allow for possible serial correlation in the regression residuals. The regression results are reported in Table 3.

The regressions provide strong evidence that the violations of the martingale restriction are related to market frictions. In particular, the average bid-ask spread of the options used in estimating the implied index value is positive and highly significant in both regressions. This means that the cost of taking a synthetic position in the index via the options market increases with the cost of trading options. This is intuitive since taking a synthetic position in the index requires an investor to incur these higher transaction costs.

Similarly, the results indicate that the implied cost of the index in the options market decreases as the liquidity of the options market improves. The coefficient for the total open interest is negative and significant in all of the regressions reported. As the total open interest increases, the options market becomes more liquid, and the cost of a synthetic position in the index becomes less expensive. In addition, the results suggest that as the level of trading activity increases, the cost of a synthetic position in the index decreases. In the first regression, trading volume is significant. In the other regressions, which include the current and lagged index returns, the coefficient for trading volume is negative but not always significant.

The second and third regressions examine whether there is evidence of path-dependent effects on option pricing. As shown, there is a pronounced negative relation between the percentage pricing differences and index returns. The negative relation is strongest for the contemporaneous index return, but is still highly significant for the returns for the previous two trading days. Intuitively, this negative relation suggests that the cost of the index in the options market increases when the market is declining or has recently declined.

Finally, the regressions indicate that time to expiration, moneyness, and the volatility proxy are all related to the violations of the martingale restriction. The implied cost of the equity component of the options increases with the average time to expiration, decreases as the average moneyness of the calls increases, and increases when the volatility of the market goes up. The relation between the martingale

[8] In contrast, Manaster and Rendleman (1982) regress stock returns on the percentage pricing differences for individual stocks. They find some evidence that the pricing differences have explanatory power for stock returns.

Table 3
Results from regressing percentage differences between the implied index value and the actual index value on the indicated variables

Int	*M*	*T*	*A Ret*	$A\,Ret_{-1}$	$A\,Ret_{-2}$	*BA*	*OI*	*CV*	*N*	*Ret*	Ret_{-1}	Ret_{-2}	R^2
Coefficient													
0.00441	−0.00031	0.01421	0.06964	0.04570	0.02429	0.01887	−6.534	−6.039	−0.00026				0.644
0.00433	−0.00027	0.01653	0.06574	0.04380	0.02351	0.01648	−6.168	−5.205	−0.00024	−0.02814			0.652
0.00464	−0.00024	0.01787	0.06012	0.04147	0.02270	0.01458	−6.811	−4.543	−0.00022	−0.04027	−0.02338	−0.03176	0.663
t-statistic													
4.74	−8.12	2.67	4.33	2.83	1.69	5.20	−2.83	−2.00	−8.25				
4.66	−6.69	3.13	4.13	2.74	1.65	4.55	−2.66	−1.74	−7.47	−3.20			
5.01	−6.10	3.42	3.82	2.62	1.61	4.05	−2.94	−1.54	−7.10	−4.16	−2.39	−3.49	

The table reports the coefficients and *t*-statistics for each of the indicated independent variables. *Int* is the regression intercept, *M* is the average moneyness of the options, *T* is the average time to expiration, *A Ret*, $A\,Ret_{-1}$, and $A\,Ret_{-2}$ are the current and first two lagged absolute daily returns on the S&P 100 index, *BA* is the average bid-ask spread of the options, *OI* is open interest, *CV* is the total trading volume for all calls for that day, *N* is the number of calls used to compute the implied index value for that day, and *Ret*, Ret_{-1}, and Ret_{-2} are the current and first two lagged daily returns on the S&P 100. The coefficients for *OI* and *CV* are multiplied by 10^9. The sample consists of 444 daily observations.

restriction and the biases of the Black-Scholes model are examined in more depth in the following section.

5. Relaxing the Martingale Restriction

Because the Black-Scholes option pricing model is a no-arbitrage model, it imposes the martingale restriction on the mean of the lognormal risk-neutral density. As shown, this martingale restriction is strongly rejected by the data. In this section, I examine whether relaxing the martingale restriction in the Black-Scholes model improves the performance of the model in describing actual option prices.

I designate the traditional version of the Black-Scholes model the restricted model since it imposes the martingale restriction. Recall that the restricted model implies that the risk-neutral density is lognormal, where the mean of the lognormal is fully specified by the martingale restriction. Now consider a version of the Black-Scholes model in which the pricing density is lognormal, but the martingale restriction is not imposed. I term this model the unrestricted Black-Scholes model. Intuitively, the unrestricted model can be viewed as a simplistic 'general equilibrium' form of the Black-Scholes model.

I compare the performance of the restricted and unrestricted Black-Scholes models in the following way. First, I estimate the implied volatility that best fits the call prices for each day in the sample. For example, if there are 15 options in the sample for day T, I estimate the implied volatility which minimizes the sum of squared pricing errors for the 15 options. I then repeat this procedure for each of the 444 days in the sample period, resulting in 444 daily implied volatility estimates and 4582 pricing errors (444 days times an average of 10.32 calls). The pricing errors are computed as the difference between the actual call price and the price implied by the fitted model. This gives the implied volatility estimates and pricing errors for the restricted model. I estimate the implied volatility and pricing errors for the unrestricted model in a similar fashion. Rather than estimating only the implied volatility, however, the unrestricted model is estimated by jointly finding the implied index value and volatility that best fits the option prices. Note that both the pricing errors and implied volatility estimates of the restricted model will differ from those for the unrestricted model.

Table 4 compares the pricing errors for the restricted model with those of the unrestricted model. Because the unrestricted model has an extra parameter, the pricing errors of the unrestricted model should be less than the restricted model. Surprisingly, however, the pricing errors from the unrestricted model are dramatically less than those of the restricted model. In particular, the median absolute pricing error for

Table 4
Summary statistics for the absolute value of the differences between the fitted and actual call prices

	Mean	Std. dev.	Min.	Median	Max.
Restricted model	0.3801	0.2918	0.0001	0.3203	2.0975
Unrestricted model	0.1783	0.1679	0.0000	0.1333	1.4888

The restricted model is the standard Black-Scholes model and is fitted each day by implying a single implied volatility estimate for all of the call options available for that day. The unrestricted model is the Black-Scholes model fitted each day by implying a single index value and volatility estimate for all of the call options available for that day. The total sample size is 4582.

the restricted model is 32 cents. In contrast, the median absolute pricing error for the unrestricted model is approximately 13 cents, which is about 40 percent of that for the restricted model. By relaxing the martingale restriction, most of the pricing error of the Black-Scholes model is eliminated.

I also examine the degree to which the pricing errors of each model display systematic patterns or biases. Specifically, I examine whether the pricing errors are purely random, as would be the case if they were due to measurement error, or whether they are related to option-specific characteristics as their moneyness, time to expiration, or bid-ask spread.

Figure 2 graphs the pricing errors from the restricted model against the moneyness of the options. As many studies have shown, the traditional Black-Scholes model displays a significant strike price bias. In contrast, the pricing errors for the unrestricted model are graphed in Figure 3. As shown, the unrestricted model results in significantly less bias than the restricted model.

Figure 4 graphs the pricing errors from the restricted model against the time to expiration of the options. Figure 5 shows the same graph for the pricing errors from the unrestricted model. There is a clear negative time to expiration bias in the restricted model. In contrast, the pricing errors from the unrestricted model display a slightly positive time to expiration bias. Note, however, that the magnitude of this bias is much smaller than in the restricted model.

The pricing errors from the restricted model are plotted against the bid-ask spread of the options in Figure 6. This graph indicates that the traditional Black-Scholes model has a strong bid-ask spread bias. The Black-Scholes model tends to overprice options with a small bid-ask spread and underprice options with a large bid-ask spread. In contrast, Figure 7 shows that most of the bid-ask spread bias in the pricing errors is eliminated in the unrestricted model.

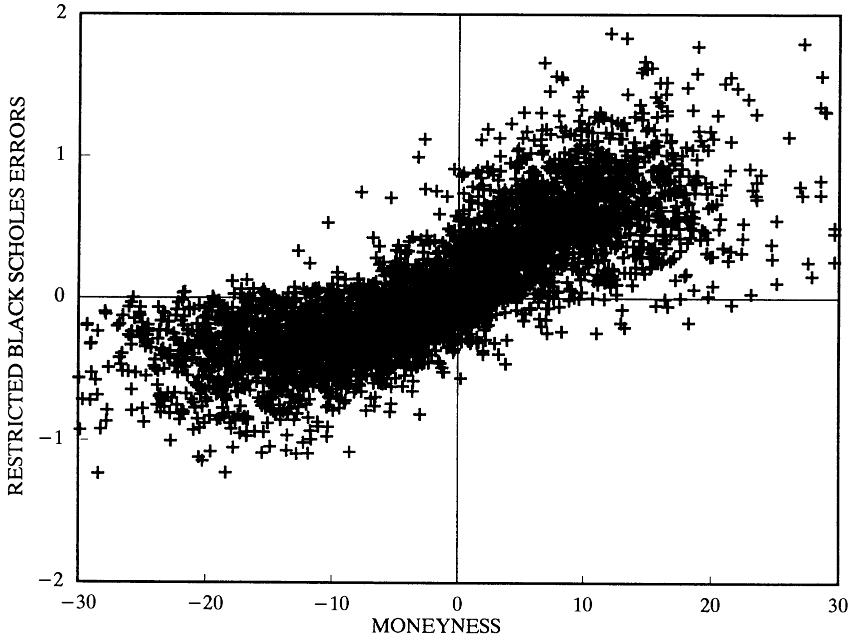

Figure 2
Pricing errors from the restricted Black-Scholes model graphed against the moneyness of the options
The pricing errors are computed by solving for the Black-Scholes implied volatility estimate that results in the best fit to the S&P 100 index call option prices in the sample. The pricing errors represent the difference between the actual and fitted prices.

Table 5 reports the results of regressing the pricing errors from the restricted and unrestricted models on the moneyness, time to expiration, and bid-ask spreads of the options. I use these variables since I am examining the cross-sectional properties of the pricing errors and these variables are option specific. I do not include variables common to all options such as the volatility proxy since they do not provide explanatory power for the cross section of pricing differences. These regressions indicate that over 62 percent of the variation in the pricing differences from the restricted model is due to moneyness, time to expiration, and bid-ask spread bias. These regressions also indicate that the bid-ask spread bias is distinct from the other previously documented biases of the Black-Scholes model. In contrast, less than 21 percent of the variation in the pricing differences from the unrestricted model is due to these biases.

Since the unrestricted model does not impose the martingale restriction, the implied volatility estimates obtained from the unrestricted

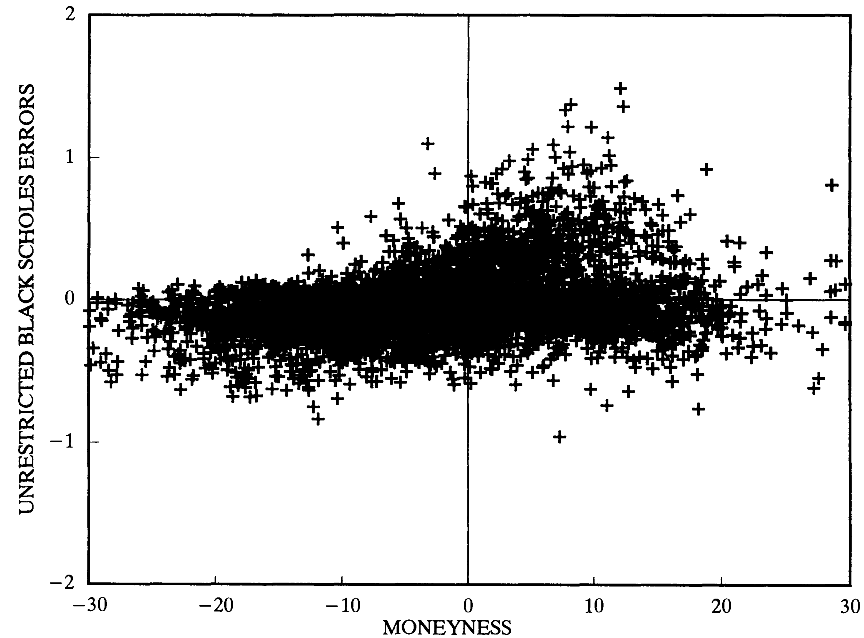

Figure 3
Pricing errors from the unrestricted Black-Scholes model graphed against the moneyness of the options
The pricing errors are computed by solving for the Black-Scholes implied index value and volatility estimates that result in the best fit to the S&P 100 index call option prices in the sample. The pricing errors represent the difference between the actual and fitted prices.

Table 5
Results from regressing differences between the fitted and actual call prices on the indicated variables

	Int	*M*	*T*	*BA*	R^2
Restricted					
Coefficient	0.03510	0.02624	−0.62431	0.87409	0.624
t-statistic	3.27	41.57	−6.14	14.33	
Unrestricted					
Coefficient	−0.17651	0.00736	1.7054	0.14667	0.209
t-statistic	−22.15	15.67	22.55	3.23	

The restricted model is the standard Black-Scholes model and is fitted each day by implying a single implied volatility estimate for all of the call options available for that day. The unrestricted model is the Black-Scholes model fitted each day by implying a single index value and volatility estimate for all of the call options available for that day. *Int* is the regression intercept, *M* is the average moneyness of the options, *T* is the average time to expiration, and *BA* is the average bid-ask spread of the option. The total sample size is 4582.

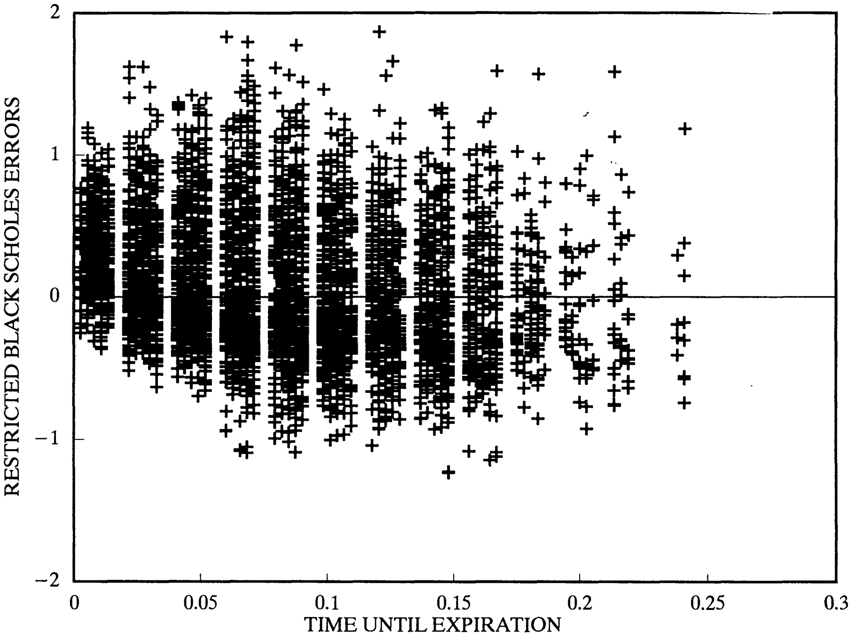

Figure 4
Pricing errors from the restricted Black-Scholes model graphed against the time until expiration of the options
The pricing errors are computed by solving for the Black-Scholes implied volatility estimate that results in the best fit to the S&P 100 index call option prices in the sample. The pricing errors represent the difference between the actual and fitted prices.

model need not equal those estimated from the restricted model. Figure 8 graphs the difference between the restricted and unrestricted implied volatility estimates for the 444 days in the sample period. As shown, the implied volatility estimates are quite different. The restricted implied volatility is nearly always higher than the unrestricted implied volatility. On average the implied volatility estimated using the restricted model is 0.1939. In contrast, the implied volatility estimated using the unrestricted model is 0.1702. Intuitively, this is because the restricted model implies a lower mean for the risk-neutral density than the unrestricted model. Thus, the implied volatility of the restricted model must be higher to compensate for imposing the martingale restriction.

An important implication of these results is that if the martingale restriction imposed by an option pricing model is rejected by the data, the implied volatility estimates obtained from the model may not be reliable estimates of the actual volatility of the underlying asset.

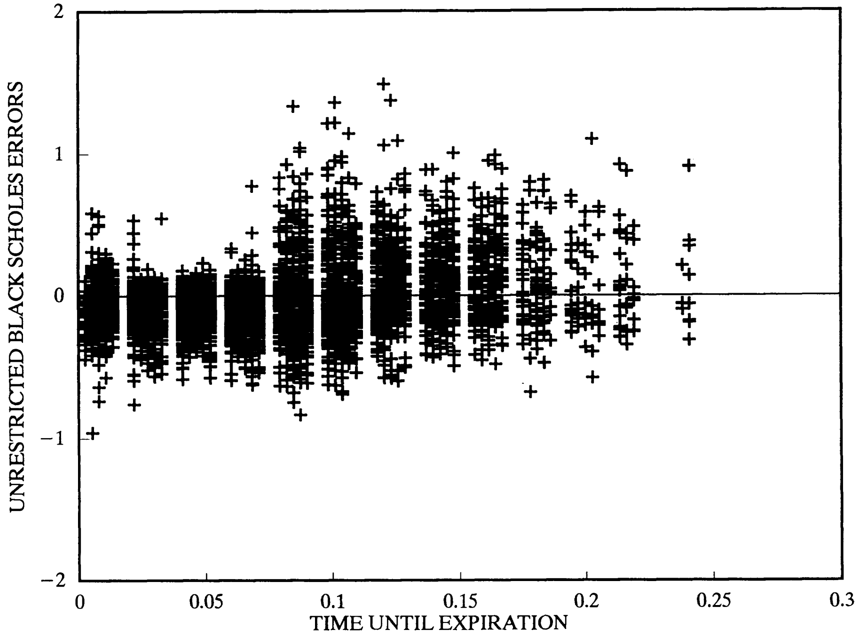

Figure 5
Pricing errors from the unrestricted Black-Scholes model graphed against the time until expiration of the options
The pricing errors are computed by solving for the Black-Scholes implied index value and volatility estimates that result in the best fit to the S&P 100 index call option prices in the sample. The pricing errors represent the difference between the actual and fitted prices.

This may provide an explanation why recent evidence by Canina and Figlewski (1993) suggests that the implied volatility estimates from the Black-Scholes model contain little information about the realized volatility of the underlying asset.

Finally, I note that even very simple modifications of the Black-Scholes model that allow for the effects of market frictions can substantially improve on its performance. For example, simply multiplying the index value by 1.004 before inputting it into the Black-Scholes formula and estimating the implied volatility and pricing errors reduces the median absolute pricing error from 32 cents to 16 cents. Similarly, the mean absolute pricing error is reduced from 38 cents to 26 cents. The R^2 in the cross-sectional regression of pricing errors on moneyness, time to expiration, and bid-ask spreads is reduced from 0.62 to 0.16. Clearly, multiplying the index value by a factor of 1.004 is an ad hoc adjustment. Nevertheless, it suggests that a fully developed general equilibrium option pricing model which explicitly incorpo-

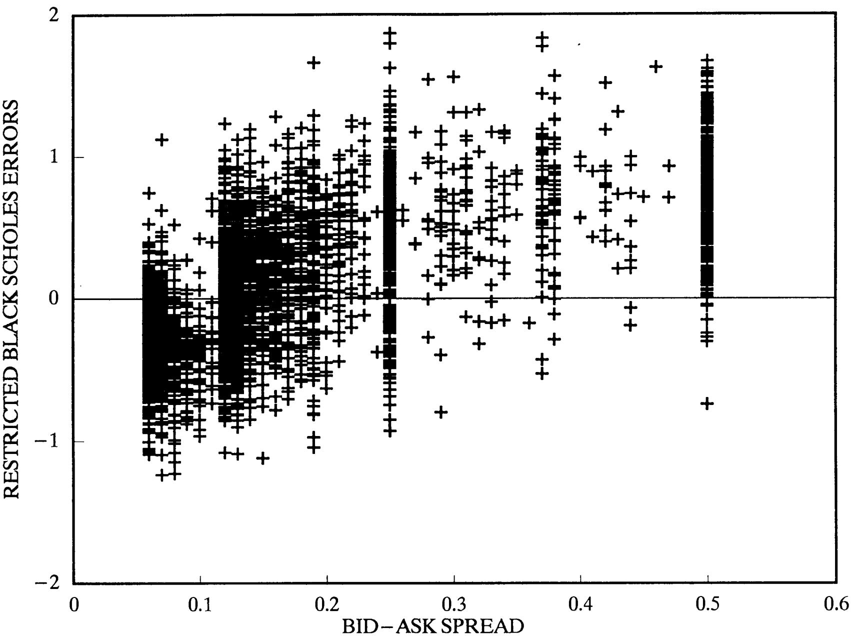

Figure 6
Pricing errors from the restricted Black-Scholes model graphed against the bid-ask spread of the options
The pricing errors are computed by solving for the Black-Scholes implied volatility estimate that results in the best fit to the S&P 100 index call option prices in the sample. The pricing errors represent the difference between the actual and fitted prices.

rates transaction costs, market illiquidity, and other market frictions would likely lead to large improvements in pricing performance over the traditional Black-Scholes model.

6. Testing Alternative Models

Earlier, I examined whether the martingale restriction was satisfied by the Black-Scholes model. In this section, I examine a general no-arbitrage option pricing model that nests or closely approximates many of the option pricing models that have appeared in the literature.

In developing this general option pricing model, I specify the functional form of the risk-neutral density rather than specifying the dynamics of the underlying asset. In doing this, our goal is to specify the risk-neutral density in a way that allows for the broadest set of possible shapes for the risk-neutral density.

Assume that the risk-neutral density is a member of the class of four-parameter density functions known as the Fourier series or Edgeworth

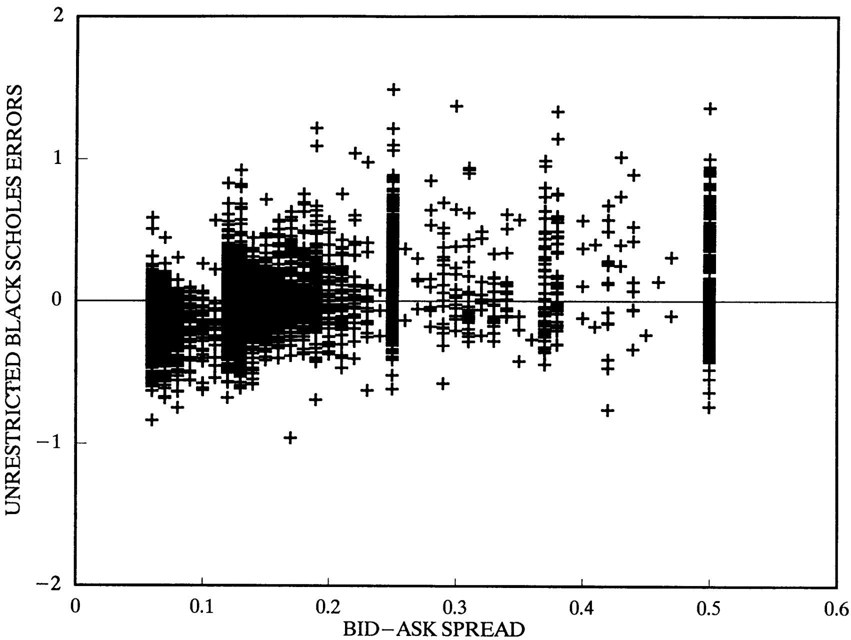

Figure 7
Pricing errors from the unrestricted Black-Scholes model graphed against the bid-ask spread of the options
The pricing errors are computed by solving for the Black-Scholes implied index value and volatility estimates that result in the best fit to the S&P 100 index call option prices in the sample. The pricing errors represent the difference between the actual and fitted prices.

expansion family of densities. As shown by Johnson and Kotz (1970), the four parameters can be chosen in a way to match the first four moments of any continuous density function. Thus, this family of density functions admits a virtually unlimited class of possible shapes. These density functions are closely related to the approximating functions used in neutral network estimation and have been used in Gallant, Hansen and Tauchen (1990), Gallant and Nychka (1987), Gallant and Tauchen (1989), Jarrow and Rudd (1982), and Singleton (1990).

Let Z be the standardized value of $\ln X_T$:

$$Z = \frac{\ln X_T - \alpha}{\sigma}, \tag{7}$$

where α and σ are the conditional mean and standard deviation of $\ln X_T$ implied by the risk-neutral pricing measure Q. Furthermore, let $q(Z)$ denote the risk-neutral density for Z implied by Q. Following

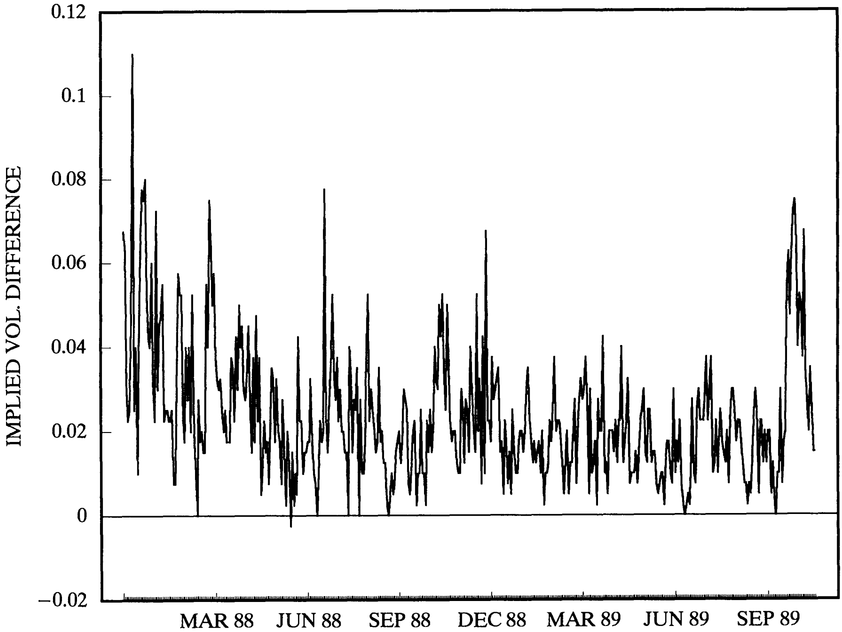

Figure 8
The implied volatility from the restricted Black-Scholes model minus the implied volatility from the unrestricted Black-Scholes model
In the restricted Black-Scholes model, only the implied volatility is estimated from option prices. In the unrestricted Black-Scholes model, both the implied index value and the implied volatility are estimated from option prices.

Johnson and Kotz (1970), let $q(Z)$ be of the form

$$q(Z) = \frac{\exp(-Z^2/2)}{\sqrt{2\pi}}(1 + \beta(Z^3 - 3Z) + \gamma(Z^4 - 6Z^2 + 3)), \tag{8}$$

where β and γ are coefficients related to the higher moments of $\ln X_T$. Together Equations (7) and (8) specify Edgeworth expansion family of risk-neutral density functions defined by the four parameters α, β, γ, and σ.

With the risk-neutral density specified, the risk-neutral valuation operator in Equation (4) can be used to express the price of a European call option on X_T with strike price K as

$$D_T \int_{-\infty}^{\infty} \max(0, \exp(\alpha + \sigma Z) - K) q(Z)\, dZ. \tag{9}$$

This integral can be evaluated directly to give a lengthy closed-form expression for the value of a call option in terms of the cumulative normal and gamma distribution functions and the four parameters α, β, γ, and σ.

This general option pricing model includes many other option pricing models as special cases. Examples of models that are nested within this general model include

- the Black-Scholes (1973) model,
- the Merton (1973) stochastic interest rate model,
- the Merton (1976, eq. 17) jump diffusion model, and
- the Merton (1976, eq. 18) jump diffusion model.

In addition, since the risk-neutral density in Equation (8) can match the first four moments of any continuous density, the four parameters of the model can be chosen in such a way that Equation (9) closely approximates most existing option pricing models.

Given market prices for four call options differing only in their strike prices, I can invert the expressions for the call values to solve for the four parameters. Since the parameters are conditional on T, I estimate the parameters separately for each different time to expiration. On average, there are 1.83 sets of four or more options with common expiration dates for each day during the sample. Thus, the total number of risk-neutral densities estimated is $1.83 \times 444 = 812$. In estimating the parameters, I use only the four call options that are closest to the money. This provides another diagnostic check on earlier results which use all available options. Once the paramters of a risk-neutral density are determined, the first moment of the risk-neutral density is given by the expression

$$E_Q[X_T] = \exp(\alpha + \sigma^2/2)(1 + \beta\sigma^3 + \gamma\sigma^4). \tag{10}$$

The martingale restriction for this general option pricing model can now be expressed as

$$X_0 = D_T \exp(\alpha + \sigma^2/2)(1 + \beta\sigma^3 + \gamma\sigma^4). \tag{11}$$

This approach to estimating the conditional mean of the risk-neutral density parallels recent work by Gallant, Hansen, and Tauchen (1990) and Hansen and Jagannathan (1991). In these papers, a semiparametric approach is applied to security market data to place bounds on the admissible region for means and standard deviations of the intertemporal marginal rate of substitution for consumers. In the context of this paper, this is equivalent to placing bounds on the moments of the distribution of ρ defined in Equation (1).

The results from the martingale tests for this model are reported in Table 6. Even though the class of possible types of density functions is much larger, the results are strikingly similar to those for the Black-

Scholes model. Table 6 shows that a large majority of the percentage differences are positive. Of the 812 estimates, 759 or 93.5 percent are greater than zero. This pattern is the same for all of the individual maturity categories. In fact, for the four-week maturity category, 100 percent of the percentage differences are greater than zero. A standard binomial test strongly rejects the hypothesis that positive and negative estimates are equally likely. The overall z-statistic for the binomial test is 24.78. The mean percentage difference fo all 812 observations is 0.400 with a t-statistic of 31.08. Note that this value is only slightly less than the mean value of the percentage differences for the Black-Scholes model given in Table 2.

These results suggest that the rejection of the martingale restriction is not limited to the Black-Scholes model and that the martingale restriction is violated by a variety of more general no-arbitrage option pricing models. It is important to acknowledge, however, that there are fewer degrees of freedom available in estimating the more general model since more parameters are estimated. Thus, the parameters may not be as precisely estimated as in earlier sections.

7. Conclusion

In this article, I examined whether the martingale restriction imposed by the no-arbitrage option pricing framework holds for S&P 100 index option prices. I found that the implied index value is nearly always higher than the actual index value. The percentage differences between the implied and actual index values are related to a number of transaction cost and option market liquidity variables. These results provide evidence of a number of previously undocumented biases in the Black-Scholes model, such as a bid-ask spread bias and an open-interest bias.

There are several possible interpretations of this evidence. For example, these results may imply that market frictions have a major effect on the pricing of options. If so, then options should be valued using equilibrium rather than no-arbitrage models. I find that a simplistic 'general equilibrium' version of the Black-Scholes model in which the martingale restriction is relaxed significantly improves on the pricing performance of the traditional Black-Scholes model. A fully developed general equilibrium model would likely lead to additional pricing improvements.

It is important to acknowledge, however, that there is an alternative interpretation for these results. It may be that the actual risk-neutral density implicit in the market's valuation of options is fundamentally different from the risk-neutral densities implied by the option pricing models that I consider. This possibility is relevant since it is always

Table 6
Summary statistics for the estimates of the percentage differences between the index value implied by the general option pricing model and the current index net of dividends

Weeks until expiration	Proportion positive	z-statistic	Mean	Standard deviation	t-statistic	Minimum	1st quartile	Median	3rd quartile	Maximum	N
1	0.961	8.03	0.214	0.173	10.75	−0.037	0.107	0.179	0.281	0.924	76
2	0.944	9.18	0.284	0.238	12.37	−0.133	0.118	0.226	0.442	1.220	107
3	0.960	9.15	0.357	0.240	14.80	−0.095	0.184	0.335	0.503	1.110	99
4	1.000	10.10	0.490	0.329	15.07	0.010	0.236	0.403	0.770	1.415	102
5	0.990	10.00	0.519	0.356	14.88	−0.091	0.261	0.403	0.727	1.683	104
6	0.978	9.07	0.482	0.380	12.04	−0.313	0.231	0.358	0.724	1.666	90
7	0.921	7.34	0.501	0.465	9.39	−0.181	0.159	0.342	0.851	2.004	76
8	0.889	6.17	0.461	0.447	8.19	−0.175	0.212	0.363	0.625	1.988	63
9	0.878	5.29	0.424	0.448	6.63	−0.379	0.128	0.322	0.576	1.569	49
≥ 10	0.609	1.47	0.170	0.435	2.65	−0.484	−0.125	0.172	0.331	1.850	46
All	0.935	24.78	0.400	0.367	31.08	−0.484	0.165	0.311	0.547	2.004	812

Proportion positive is the fraction of observations that result in a positive estimate of the pricing difference. The z-statistic tests the hypothesis that the proportion of positive estimates is 0.50, and is distributed as a standard normal variate. The t-statistic tests the hypothesis that the mean estimate is zero. N denotes the number of observations.

possible to construct an $N + 1$ parameter risk-neutral density that will exactly match N option prices with the same expiration date and satisfy the martingale restriction on the mean of the risk-neutral density. For example, this can be done using techniques similar to those presented in Banz and Miller (1978), Bick (1982), and Breeden and Litzenberger (1978).

In order for this alternative explanation to account for these results, however, the risk-neutral density implicit in the market's valuation would need to be very complex since I am able to reject a model in which the risk-neutral density can match the first four moments of any continuous density function. Furthermore, it is difficult to explain why transaction costs and market liquidity measures should be related to the violations of the martingale restriction if misspecification of the functional form of the risk-neutral density was the underlying reason for the rejection of the martingale restriction. Clearly, future research should investigate in more depth the role that market frictions play in the valuation of derivative securities.

References

Amihud, Y., and H. Mendelson, 1986, "Asset Pricing and the Bid Ask Spread," *Journal of Financial Economics*, 17, 223–249.

Banz, R. W., and M. H. Miller, 1978, "Prices for State Contingent Claims: Some Estimates and Applications," *Journal of Business*, 51, 653–672.

Bates, D. S., 1991, "The Crash of '87: Was it Expected? The Evidence from Options Markets," *Journal of Finance*, 46, 1009–1044.

Bick, A., 1982, "Comments on the Valuation of Derivative Assets," *Journal of Financial Economics*, 10, 331–345.

Black, F., 1975, "Fact and Fantasy in the Use of Options," *Financial Analysts Journal*, 31, 899–908.

Black, F., and M. Scholes, 1973, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81, 637–659.

Boyle, P. P., and T. Vorst, 1992, "Option Replication in Discrete Time with Transaction Costs," *Journal of Finance*, 47, 271–293.

Breeden, D. T., 1979, "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities," *Journal of Financial Economics*, 7, 265–296.

Breeden, D. T., and R. H. Litzenberger, 1978, "Prices of State Contingent Claims Implicit in Option Prices," *Journal of Business*, 51, 621–651.

Brenner, M., G. Courtadon, and M. Subrahmanyam, 1987, "The Valuation of Stock Index Options," working paper, New York University.

Canina, L., and S. Figlewski, 1993, "The Informational Content of Implied Volatility," *Review of Financial Studies*, 6, 659–681.

Chiras, D. P., and S. Manaster, 1978, "The Informational Content of Option Prices and a Test of Market Efficiency," *Journal of Financial Economics*, 5, 213–234.

Constantinides, G. M., 1986, "Capital Market Equilibrium with Transaction Costs," *Journal of Political Economy*, 94, 842–862.

Constantinides, G. M., 1989, "Theory of Valuation: Overview and Recent Developments," S. Bhattacharya and G. M. Constantinides (eds.), in *Theory of Valuation, Frontiers of Modern Financial Theory, Vol. 1*, Rowman & Littlefield, Totowa, N.J.

Constantinides, G. M., 1993, "Option Pricing Bounds with Transaction Costs," working paper, The University of Chicago.

Cox, J. C., J. E. Ingersoll, and S. A. Ross, 1985a, "An Intertemporal General Equilibrium Model of Asset Prices," *Econometrica*, 53, 363–384.

Cox, J. C., J. E. Ingersoll, and S. A. Ross, 1985b, "A Theory of the Term Structure of Interest Rates," *Econometrica*, 53, 385–407.

Cox, J. C., and S. A. Ross, 1976, "The Valuation of Options for Alternative Stochastic Processes," *Journal of Financial Economics*, 3, 145–166.

Fackler, P. F., 1986, The Informational Content of Option Premiums, unpublished PhD dissertation, University of Minnesota.

Fackler, P. F., and R. P. King, 1990, "Calibration of Option-Based Probability Assessments in Agriculture Commodity Markets," *American Journal of Agricultural Economics*, 72, 73–83.

Gallant, A. R., L. P. Hansen, and G. Tauchen, 1990, "Using Conditional Moments of Asset Payoffs to Infer the Volatility of Intertemporal Marginal Rates of Substitution," *Journal of Econometrics*, 45, 141–179.

Gallant, A. R., and D. W. Nychka, 1987, "Seminonparametric Maximum Likelihood Estimation," *Econometrica*, 55, 363–390.

Gallant, A. R., and G. Tauchen, 1989, "Seminonparametric Estimation of Conditionally Constrained Heterogeneous Processes: Asset Pricing Applications," *Econometrica*, 57, 1091–1120.

George, T. J., and F. A. Longstaff, 1993, "Bid-Ask Spreads and Trading Activity in the S&P 100 Index Options Market," *Journal of Financial and Quantitative Analysis*, 28, 381–397.

Gibson, R., 1991, *Option Valuation*, McGraw Hill, New York.

Grundy, B. D., 1991, "Option Prices and the Underlying Asset's Return Distribution," *Journal of Finance*, 46, 1045–1070.

Hansen, L. P., and R. Jagannathan, 1991, "Implications of Security Market Data for Models of Dynamic Economies," *Journal of Political Economy*, 99, 225–262.

Hansen, L. P., and S. F. Richard, 1987, "The Role of Conditioning Information in Deducing Testable Restrictions Implied by Dynamic Asset Pricing Models," *Econometrica*, 55, 587–613.

Hansen, L. P., and K. J. Singleton, 1983, "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Stock Market Returns," *Journal of Political Economy*, 91, 249–265.

Harrison, J. M., and D. M. Kreps, 1979, "Martingales and Arbitrage in Multiperiod Securities Markets," *Journal of Economic Theory*, 20, 381–408.

Harrison, J. M., and S. R. Pliska, 1981, "Martingales and Stochastic Integrals in the Theory of Continuous Trading," *Stochastic Processes and their Applications*, 11, 261–271.

Harvey, C. R., and R. E. Whaley, 1992, "Dividends and S&P 100 Index Option Valuation," *Journal of Futures Markets*, 12, 123–137.

Heath, D., R. Jarrow, and A. Morton, 1992, "Bond Pricing and the Term Structure of Interst Rates: A New Methodology for Contingent Claims Valuation," *Econometrica*, 60, 77–105.

Hutchinson, J., A. W. Lo, and T. Poggio, 1994, "A Nonparametric Approach to Pricing and Hedging Derivative Securities Via Learning Networks," *The Journal of Finance 49*, 851–889.

Jarrow, R., and A. Rudd, 1982, "Approximate Option Valuation for Arbitrary Stochastic Processes," *Journal of Financial Economics*, 10, 347–369.

Johnson, N. L., and S. Kotz, 1970, *Continuous Univariate Distributions—1 and 2*, John Wiley & Sons, New York.

Latane, H., and R. Rendleman, Jr., 1976, "Standard Deviations of Stock Price Ratios Implied in Options Prices," *Journal of Finance*, 31, 369–381.

Leland, H., 1985, "Option Pricing and Replication with Transaction Costs," *Journal of Finance*, 40, 1283–1301.

Levy, H., 1985, "Upper and Lower Bounds of Put and Call Option Value: Stochastic Dominance Approach," *Journal of Finance*, 40, 1197–1217.

Lo, A. W., 1987, "Semi-Parametric Upper Bounds for Option Prices and Expected Payoffs," *Journal of Financial Economics*, 19, 373–387.

Lucas, R. E., 1978, "Asset Prices in an Exchange Economy," *Econometrica*, 46, 1429–1145.

Macbeth, J. D., and L. J. Merville, 1980, "Tests of the Black-Scholes and Cox Call Option Valuation Models," *Journal of Finance*, 35, 285–300.

Madan, D. B., and F. Milne, 1994, "Contingent Claims Valued and Hedged by Pricing and Investing in a Basis," *Mathematical Finance*, 4, 223–245.

Manaster, S., and R. J. Rendleman, 1982, "Option Prices as Predictors of Equilibrium Stock Prices," *Journal of Finance*, 37, 1043–1057.

Merton, R. C., 1973, "Theory of Rational Option Pricing," *Bell Journal of Economics and Management Science*, 4, 141–183.

Merton, R. C., 1976, "Option Pricing When Underlying Stock Returns are Discontinuous," *Journal of Financial Economics*, 3, 125–144.

Perrakis, S., and P. J. Ryan, 1984, "Option Pricing Bounds in Discrete Time," *Journal of Finance*, 39, 519–525.

Phillips, S., and C. Smith, 1980, "Trading Costs for Listed Options: The Implications for Market Efficiency," *Journal of Financial Economics*, 8, 179–201.

Ritchken, P., 1985, "On Option Pricing Bounds," *Journal of Finance*, 40, 1219–1233.

Rubinstein, M., 1976, "The Valuation of Uncertain Income Streams and the Pricing of Options," *Bell Journal of Economics*, 7, 407–425.

Rubinstein, M., 1985, "Nonparametric Tests of Alternative Option Pricing Models Using All Reported Trades and Quotes on the 30 Most Active CBOE Option Classes from August 23, 1976 Through August 31, 1978," *Journal of Finance*, 40, 455–480.

Rudin, W., 1987, *Real and Complex Analysis, 3rd ed.*, McGraw-Hill, New York.

Schmalensee, R., and R. R. Trippi, 1978, "Common Stock Volatility Expectations Implied by Option Prices," *Journal of Finance*, 33, 129–147.

Sherrick, B. J., S. H. Irwin, and D. L. Forster, 1990, "Expected Soybean Futures Price Distributions: Option-Based Assessments," *The Review of Futures Markets*, 9, 386–409.

Sherrick, B. J., S. H. Irwin, and D. L. Forster, 1992, "Option Based Evidence of the Nonstationarity of Expected S&P 500 Future Price Distribution," *Journal of Futures Markets*, 12, 275–290.

Shimko, D. C., 1993, "Beyond Implied Volatility: Probability Distributions and Hedge Ratios Implied by Option Prices," *Risk Magazine* 6(4), 33–37.

Tuckman, B., and J. L. Vila, 1992, "Holding Costs and Equilibrium Arbitrage," working paper, New York University.

Vayanos, D., and J. L. Vila, 1992, "Equilibrium Interest Rate and Liquidity Premia under Proportional Transaction Costs," working paper, Massachussetts Institute of Technology.