Pricing S&P 500 Index Options Using a Hilbert Space Basis

Peter A. Abken, Dilip B. Madan, and
Sailesh Ramamurtie

Federal Reserve Bank of Atlanta
Working Paper 96-21
December 1996

Abstract: This paper tests the approach of Madan and Milne (1994) and its extension in Abken, Madan, and Ramamurtie (1996) for pricing contingent claims as elements of a separable Hilbert space. We specialize the Hilbert space basis to the family of Hermite polynomials and test the model on S&P 500 index options. Restrictions on the prices of Hermite polynomial risk are imposed that allow all option maturity classes to be used in estimation. These restrictions are rejected by our empirical tests of a four-parameter specification of the model. Nevertheless, the unrestricted four-parameter model, based on a single maturity class, demonstrates better out-of-sample performance than that of the Black-Scholes version of the Hermite model. The unrestricted four-parameter model results indicate skewness and excess kurtosis in the implied risk-neutral density. The skewness of the risk-neutral density contrasts with the symmetry of the statistical density estimated using the Hermite model on the S&P 500 index returns.

JEL classification: G13, C52

Please address questions of substance to Peter A. Abken, Senior Economist, Research Department, Federal Reserve Bank of Atlanta, 104 Marietta Street, N.W., Atlanta, Georgia 30303-2713, 404/521-8783, 404/521-8810 (fax), peter.abken@atl.frb.org; Dilip B. Madan, College of Business and Management, University of Maryland, College Park, Maryland 20742, 301/405-2127, 301/314-9157 (fax), dilip@bmgtmail.umd.edu; and Sailesh Ramamurtie, Department of Finance, College of Business Administration, Georgia State University, Atlanta, Georgia 30303, 404/651-2710, 404/651-2630 (fax), sramamurtie@gsu.edu.

Questions regarding subscriptions to the Federal Reserve Bank of Atlanta working paper series should be addressed to the Public Affairs Department, Federal Reserve Bank of Atlanta, 104 Marietta Street, N.W., Atlanta, Georgia 30303-2713, 404/521-8020, http://www.frbatlanta.org.

Pricing S&P 500 Index Options Using a Hilbert Space Basis

This paper tests the option pricing approach of Madan and Milne (1994) and its extension in Abken, Madan, and Ramamurtie (1996) on a sample of S&P 500 index options, which have served as a testing ground for option pricing models. Formally, the Madan and Milne model prices contingent claims as elements of a separable Hilbert space. The Black-Scholes model is a parametric special case of the Madan and Milne model. Pricing in terms of a Hilbert space basis is analogous to the use of discount bonds as a basis for pricing fixed income securities or the construction of branches of a binomial tree in pricing options. The Hilbert space model is a semi-non-parametric approach to pricing that readily captures important features of the option data, such as the volatility smile.

The application of Madan and Milne's approach used here specializes the Hilbert space basis to the family of Hermite polynomials. Using this method, we infer the underlying risk-neutral density from traded security prices. This density could then be used to price and hedge a variety of other contingent claims. There has been considerable interest recently in estimating implied distributions from option prices. Examples include Longstaff (1992, 1995), Shimko (1993), Dupire (1994), Rubinstein (1994), Derman and Kani (1994), Jackwerth and Rubinstein (1995), Aït-Sahalia and Lo (1995), and Dumas, Fleming, and Whaley (1996). We also apply the Hilbert space basis approach to estimating the statistical density for the same sample period and compare it with the estimated risk neutral density. Our approach is readily implemented and does not require large amounts of data or computing power.

The original model of Madan and Milne could only be applied to one option maturity class at a time. The extension in Abken, Madan, and Ramamurtie (1996) enables all traded option maturities classes to be used in estimation. The extension specifies a simple restriction on the coefficients weighting the Hermite polynomial basis elements that gives an internally consistent model for pricing different maturity options.

We first estimate the Hermite basis pricing model and a Black-Scholes version using all S&P 500 index (SPX) option maturity classes (hereafter such models are referred to as restricted models). For our 1990–1992 sample period, we reject the Black-Scholes version, based on its in-sample fit, in favor of a four-parameter Hermite model. This is not surprising since we have three extra degrees of freedom to fit in sample. However, the four-parameter restricted model is rejected when tested against a four-parameter model estimated separately for a single option maturity class (the unrestricted model). The apparent cause of the rejection is an upward slope of the term structure of volatility. The model assumes a flat volatility structure. The problem was much more acute in the application of the Hilbert space basis approach to Eurodollar futures options in our previous paper, but even for index equity options, estimated volatility turns out to make a significant departure from the model's underlying assumptions.

The unrestricted results indicate skewness and excess kurtosis in the implied risk-neutral density. The statistical density estimated from a time series on the S&P 500 index also exhibits excess kurtosis of a similar magnitude to the risk-neutral density's. On the other hand, in contrast to the risk-neutral density, the estimated statistical density is symmetric.

We also examine one-day-ahead prediction errors of the restricted four-parameter model and the Black-Scholes version. The parameters of both models are estimated from a panel of options drawn from one day and used for pricing a panel of options the next trading day. The four-parameter model performs poorly compared with Black-Scholes. This confirms the outcomes of formal Wald tests of the models' coefficients, estimated for each day in the sample using the restricted and unrestricted models. In contrast, the unrestricted four-parameter model (fit to a single option maturity class) outperforms the Black-Scholes version of the model applied to the same data set. We offer an interpretation of this outcome that suggests that a more general version of the Hilbert space basis pricing model, derived in Madan and Milne (1994), could accommodate a non-constant term structure of volatility.

The paper is organized as follows: Section 1 reviews the framework of Madan and Milne (1994) and gives an overview of the density measures to be estimated. This section gives an exposition of option pricing using a Hermite polynomial basis and provides an explicit representation of a four-parameter risk-neutral density as well as that of the corresponding statistical density. Section 2 presents the application of the model to the S&P 500 index options and reviews the estimation approach. Section 3 summarizes and interprets the results for both the risk-neutral and statistical densities. The fit of the model is evaluated both in- and out-of-sample. Section 4 offers concluding remarks.

1. Option Pricing Using a Hermite Polynomial Basis

In this section we describe the model setup and assumptions underlying the Hermite polynomial approximation approach. Madan and Milne (1994) develop a model for valuation of contingent claims and static hedging strategies by identifying a set of "basis" claims and pricing them. The model has the following assumptions. There is an underlying probability space (Ω, F, P) for time $t \in [0,T]$ with a complete, increasing, and right-continuous filtration $\{F_t \,|\, 0 \le t \le T\}$, generated by a d-dimensional Brownian motion $z(t)$ initialized at zero. There exists a finite set of d primary securities that can be traded continuously. There also exists a risk-free or money market account that grows in value at an instantaneous rate given by a positive process, $r(t)$. In the theoretical economy assets can be traded continuously. The model shifts to a discrete-time setting for pricing and hedging a wide class of contingent claims, which are expressed as functions of the primary assets at specific discrete points in time. This is the first step in the discretization of the model that makes static hedging strategies feasible.

The next major assumption of the model is that the set of all contingent claims is rich enough to form a Hilbert space that is separable and for which an orthonormal basis exists as a consequence. The markets are assumed to be complete and free of arbitrage opportunities. This has the same flavor as valuation in continuous time, in which contingent claims are redundant.

However, in discrete time, the set of contingent claims forming the basis for the Hilbert space has a function akin to that of the set of Arrow-Debreu securities. The set of contingent claims forms a minimal statically spanning collection of claims, which can be used to value and statically hedge any other contingent claim. Consequently, dynamic portfolio rebalancing is unnecessary.

A drawback of this approach is that the dimensionality of the orthonormal basis set is typically very large and in general does not give any operational advantage in describing or pricing the basis. However, we assume that a finite basis can be constructed and used for pricing and hedging strategies. Once the basis representation of claims and the claim representation of basis elements are obtained, valuation and hedging—static or dynamic—of any other claim are feasible.

For practical applications of the approach, the model is restricted to a finite set of embedded discrete-time equivalents of the underlying continuous-time stochastic processes driving the money market and primary assets. Thus, contingent claims are functions of finitely many variables, and the Hilbert space is consequently separable with a finite basis.

The next major step in the application of the model lies in the change of measures. We defined above the existence of the probability space (Ω, F, P). However, an effective basis cannot be constructed without a complete knowledge of P, which is typically unknown to the practitioner. Madan and Milne rely on an approach developed by Elliot (1993), in which there is a change of measure from P to a reference measure R. The latter is assumed to be Gaussian in a discrete context. This change of measure introduces errors or deviations that are assumed to be sufficiently well bounded.

Finally, while we assume the structure of the reference measure R, what we have in practice, through discrete-time observation of the evolution of the prices (or returns) of both the primary assets and the contingent claims, is a statistical discrete-time model $(\widetilde{\Omega}, \widetilde{F}, \widetilde{R})$ that is assumed to be a sufficiently close approximation of (Ω, F, R). With all of these assumptions and approximations in hand, we proceed with the exercise of identifying the "basis" set and valuing the various contingent claims.

In the application of the model, we assume that there is only one primary asset and that the Hilbert space may be viewed as the space of functions defined on $\mathbf{R}^M$, where M represents the number of sample observations at discrete times over the interval $[0,T]$. The probability measure R is assumed to be Gaussian with a Hermite polynomial basis (see Rozanov 1982).

The statistical density $P(z)$ or risk-neutral density $Q(z)$ can be represented as a product of a change of measure density and a reference measure density:

$$\begin{aligned} P(z) &= \upsilon(z)n(z), \\ Q(z) &= \lambda(z)n(z), \end{aligned} \tag{1}$$

where $\upsilon(z)$ and $\lambda(z)$ are the statistical and risk-neutral change of measure densities, respectively. The reference measure density, $R(z) = n(z) \equiv \left(1/\sqrt{2\pi}\right)e^{-z^2/2}$, is the standard Gaussian density, where z is a standardized normal random variate. The focus of the analysis of this paper is on the specification and estimation of the change of measure densities. The change of measure densities must be bounded above and below by constants for the Hilbert space to be identical with respect to both measures. The analysis applies in any case to claims that are reference-measure square integrable. Given this assumption and that of Gaussian random process generating uncertainty, a basis for the Gaussian reference space may be constructed using Hermite polynomials. These polynomials are defined in terms of the normal density as

$$\Phi_k(z) = (-1)^k \frac{\partial^k n(z)}{\partial z^k} \frac{1}{n(z)}, \tag{2}$$

and normalized to unit variance

$$\phi_k(z) = \Phi_k(z) / \sqrt{k!}. \tag{3}$$

The Hermite polynomials form an orthonormal system.

Under the reference measure, the asset price is assumed to evolve as geometric Brownian motion:

$$S_t = S_0 e^{\mu t + \sigma\sqrt{t} z_t - \sigma^2 t/2} = S_0 e^{\eta t}, \tag{4}$$

with drift rate μ, variance rate σ, and $z_t \sim N(0,1)$. All parameters are assumed to be constant. The risk-free rate is also assumed to be constant. The exponent in equation (4) is the continuously compounded return process

$$\eta \sim N\left(\mu - \frac{\sigma^2}{2}, \frac{\sigma^2}{t}\right). \tag{5}$$

This process will be applied below for the S&P 500 index.

Madan and Milne show that any contingent claim payoff $g(z)$ can be represented in terms of the Hermite polynomial basis as

$$g(z) = \sum_{k=0}^{\infty} a_k \phi_k(z), \tag{6}$$

where

$$a_k = \int_{-\infty}^{\infty} g(z)\phi_k(z)n(z)\,dz.$$

The a_k coefficient is the covariance of the k-th Hermite polynomial risk with the contingent claim payoff, and may be interpreted as the number of units of the k-th basis-element contingent claim $\phi_k(z)$ to hold in a portfolio that replicates the contingent claim payoff. The price of this contingent claim is expressed as

$$V[g(z)] = \sum_{k=0}^{\infty} a_k \pi_k, \tag{7}$$

where π_k is the implicit price of Hermite polynomial risk $\phi_k(z)$. The market price of risk π_k may be interpreted as the forward price, for delivery upon the contingent claim's maturity, of basis-element contingent claim $\phi_k(z)$.

Equation (7) is the primary focus for our empirical work, and its empirical counterpart is discussed below. Given the structure of the contingent claims and the assumed probability model, the Hermite polynomial coefficients a_k are well defined, and hence the π_k's can be inferred from the observed prices. The parameters of the risk-neutral change of density measure $\lambda(z)$ and those of the reference measure are estimated from SPX option prices. In a separate procedure, the parameters of the statistical change of density measure $\nu(z)$ and reference measure are estimated from the underlying return process using a time-series on the S&P 500 index.

The payoff function $g(z)$ is specialized to standard European call and put payoffs in the empirical work below. The call option payoff is denoted by

$$c(z,S_0,x,\mu,\sigma,t)=[S_0 e^{\mu+\sigma\sqrt{t}z-\sigma^2 t/2}-x]^+ \tag{8}$$

and the put option payoff by

$$p(z,S_0,x,\mu,\sigma,t)=[x-S_0 e^{\mu+\sigma\sqrt{t}z-\sigma^2 t/2}]^+. \tag{9}$$

The call option payoff can be expressed in terms of the Hermite polynomial basis using a call option generating function

$$\Phi(u,S_0,x,\mu,\sigma,t)=\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty} c(z,S_0,x,\mu,\sigma,t)e^{-(z-u)^2/2}dz, \tag{10}$$

where u is a dummy variable in the generating function for Hermite polynomials:

$$\frac{e^{-(z-u)^2/2}}{\sqrt{2\pi}}=\sum_{k=0}^{\infty}\phi_k(z)n(z)\frac{u^k}{\sqrt{k!}}. \tag{11}$$

The a_k coefficients in (6) are defined as follows for the call option payoffs:

$$a_k(S_0,x,\mu,\sigma,t) = \frac{\partial^k \Phi(u,S_0,x,\mu,\sigma,t)}{\partial u^k}\Big|_{u=0} \frac{1}{\sqrt{k!}}. \tag{12}$$

Hence Φ is a call option generating function.

Upon evaluating the integral in (10), the call option generating function has the following form, as reported in Madan and Milne (1994):

$$\Phi(u,S_0,x,\mu,\sigma,t) = S_0 e^{\mu+\sigma\sqrt{t}u} N(d_1(u)) - x\, N(d_2(u)), \tag{13}$$

where

$$d_1(u) = \frac{1}{\sigma\sqrt{t}} \ln\frac{S_0}{x} + \left(\frac{\mu}{\sigma}+\frac{\sigma}{2}\right)\sqrt{t} + u,$$

and

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

The corresponding put option generating function is

$$\Psi(u,S_0,x,\mu,\sigma,t) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} p(z,S_0,x,\mu,\sigma,t) e^{-(z-u)^2/2} dz. \tag{14}$$

Put options have Hermite polynomial coefficients

$$b_k(S_0,x,\mu,\sigma,t) = \frac{\partial^k \Psi(u,S_0,x,\mu,\sigma,t)}{\partial u^k}\Big|_{u=0} \frac{1}{\sqrt{k!}}. \tag{15}$$

The empirical counterpart of equation (7), the option pricing equation in terms of the Hermite polynomial basis, is estimated below. As shown in Madan and Milne, setting the drift μ equal to the risk-free rate specializes the reference measure to the equivalent martingale measure under Black-Scholes, and the Hermite polynomial pricing model collapses to the Black-Scholes model under the parametric restriction yielding $\lambda(z) \equiv 1$.[1]

1.1 The Risk-Neutral Density

The risk-neutral density is given by

$$Q(z) = \frac{1}{\sqrt{2\pi}} \sum_{k=0}^{\infty} e^{rt} \pi_k \phi_k(z) e^{-z^2/2} \tag{16}$$

with respect to the standard normal reference measure on z. In turn, the variate z may be expressed in terms of the return process η as

$$z = \frac{\eta - (\mu - \sigma^2/2)}{\sigma/\sqrt{t}}\ . \tag{17}$$

Under the reference measure, z is normally distributed with the standard moments:

$$z \sim N(0,1);\ \mathrm{E}z = 0,\ \mathrm{E}z^2 = 1\ ,\ \mathrm{E}z^3 = 0,\ \mathrm{E}z^4 = 3, \ldots \tag{18}$$

However, the actual risk-neutral distribution of z may be non-normal; thus, the more general density specification in equation (16). For the empirical work below, the Hermite polynomial expansion of the equivalent martingale measure is truncated at the fourth order.

Using the definitions of the Hermite polynomials, the truncated density is explicitly represented as

$$\begin{aligned} Q(z) = \frac{1}{\sqrt{2\pi}} e^{-z^2/2} \Bigg[&\left(\beta_0 - \frac{\beta_2}{\sqrt{2}} + \frac{3\beta_4}{\sqrt{24}}\right) + \left(\beta_1 - \frac{3\beta_3}{\sqrt{6}}\right) z \\ &+ \left(\frac{\beta_2}{\sqrt{2}} - \frac{6\beta_4}{\sqrt{24}}\right) z^2 + \frac{\beta_3}{\sqrt{6}} z^3 + \frac{\beta_4}{\sqrt{24}} z^4 \Bigg], \end{aligned} \tag{19}$$

[1] With dividend payments on the index, the drift is set to the risk-free rate less the dividend yield, both expressed as continuously compounded rates.

where $\beta_k = e^{rt}\pi_k$, the future value of the k-th price of risk coefficient.

1.2 Restrictions on the Market Prices of Risk across Time

The market prices of Hermite polynomial risk are functionally related across time. The restriction has the following form:

$$\pi_k(s) = \pi_k(t)e^{r(t-s)}(s/t)^{k/2}, \tag{20}$$

where $s < t$. Appendix B in Abken, Madan, and Ramamurtie (1996) shows the derivation of this restriction on the hypothesis that the measure change is path independent and volatility is constant. The practical implication of equation (20) is that all traded contingent claims can be used jointly in estimation of the π_k's.

1.3 The Statistical Density

The statistical density is sometimes referred to as the actual or the true probability density. It is estimated from the underlying price. The basic object for analysis is the asset return process, η, the continuously compounded return.

The fourth-order representation of the statistical density is

$$P(z) = \sum_{k=0}^{4} \alpha_k \phi_k(z) \frac{1}{\sqrt{2\pi}\sigma} \sqrt{t} e^{-\frac{z^2}{2}}, \tag{21}$$

where z is defined by equation (17). This equation simplifies to

$$P(z) = \frac{1}{\sqrt{2\pi}\sigma} \sqrt{t} e^{-\frac{z^2}{2}}, \tag{22}$$

when $\alpha_0 = 1$ and $\alpha_k = 0$, for $k = 1 \ldots 4$, i.e., if the reference measure is the actual probability measure.

2. Application to S&P 500 Index Options

2.1 Data and Institutional Background

The S&P 500 index options traded at the Chicago Board Options Exchange are regarded as among the most liquid of all exchange-traded options and have frequently been used in option pricing studies for this reason as well as the fact that they are European options. The option prices used here are the time-stamped bid-ask quotes for the period January 2, 1990 to December 31, 1992. During this period the average daily total volume was approximately 50,000 contracts.

Although transaction prices are also available, the use of bid-ask quotations is a better choice for estimation of the Hilbert space basis pricing model. The quotation data set is 20 times as large as that for transactions, with a much greater dispersion of strikes and maturities. Longstaff (1995) cites two other important reasons for preferring quote data. One is that for out-of-the-money options, the bid-ask spread can be a substantial fraction of the quoted price; consequently, transaction prices may vary between bid and ask, generating noise in the prices. The other is that transaction prices may also depend on the size of the transaction, whereas bid-ask quotations are for a standardized 10 contracts. The CBOE market makers are obligated to honor their quotes for this transaction size.

The average of bid and ask prices is used in estimation. The risk-free rate r is a continuously compounded yield computed from the Treasury bill maturing upon or immediately after the option expiration. These yields are derived from the average of bid and ask discounts reported in the *Wall Street Journal*. The dividend rate d is also expressed on a continuously compounded basis. The reported dividend rate, from Datastream International, is the average dividend payout on the S&P 500 during a quarter, divided by the closing level of the index.

The Hermite basis model is estimated for each day in the sample. Estimation is by GMM, the use of which necessitates the construction of intraday time series on option prices. A sampling procedure is used that is similar to that in Bossaerts and Hillion (1993). The first two call and two put prices during five-minute intervals are sampled throughout the trading day, which runs from 8:30 a.m. to 3:15 p.m. CST. If fewer than two calls or two puts are available, then that cross sectional interval was dropped from the intraday time series.

The quotations were filtered in five ways before being selected. First, as in Dumas, Fleming, and Whaley (1996), quotes farther than 10 percent out-of-the-money were screened out, because such options are likely to be illiquid and their quotations stale. Second, in-the-money options were eliminated in the data sets used for parameter estimation. However, the in-the-money options are included in the data sets used for the evaluation of model prediction errors. Because the focus of this study is on the tails of the implicit distribution underlying the options, having a good dispersion of strikes in the data set is crucial. The call option quotations are on average almost two percent in the money during the sample period (see Table 1), whereas they are less than one percent out of the money in the transactions data. The results are very sensitive to the degree of dispersion. In particular, the excess kurtosis of the implicit distribution, which for the SPX options is known to be positive and significant from previous studies, such as Aït-Sahalia and Lo (1995), is not significant when estimated using the Hermite basis pricing model when all strikes are included in the daily sample. With all options in the data set, most of the information about the implicit distribution is concentrated in the left-hand side. To rectify this imbalance while retaining the essentially random sampling of the options intraday, the in-the-money options were excluded from the pool of options available for sampling.

The third filter was the exclusion of options with fewer than 30 days to maturity. These short-term options carry less information about implicit distributions, as Bates (1991) argues. Another reason for omitting them is that Treasury bill yields (used in discounting cash flows) on bills with one month or less to maturity are distorted by cash management practices and are unrepresentative of other short-term money market rates (Duffee 1996). The maturity distribution of the remaining options is less skewed to the shortest-term options when the expiring contract is excluded. This may improve the power of the GMM goodness-of-fit tests, discussed below, to reject the time-to-maturity restrictions of the Hermite basis pricing model.

Staleness in the recorded S&P 500 index is a concern with the CBOE data set. Our fourth screen attempts to reduce the occurrence of stale index values in the sampled data. Dwyer, Locke and Yu (1996) estimate, taking into account transactions costs, that the cash-futures basis converges in 5 to 7 minutes when arbitrage is profitable and up to 15 minutes when it is unprofitable. Furthermore, they note that the cash and futures indexes are particularly prone to diverge from the theoretical basis during the first half hour of trading for several reasons, including delayed openings for the stocks. We follow their suggestion of excluding quotations from the first half hour of trading.

Finally, call and put options were screened using the lower arbitrage boundary conditions for calls and puts, respectively (Cox and Rubinstein 1985, pp. 129, 145). This filter was applied conservatively because measurement errors in the index, interest rate, and expected dividends can result in apparent arbitrage violations; consequently, the ask price of the option was used in computing the restriction in order to catch relatively large violations.

Table 1 gives descriptive statistics for the data samples. There are four data samples for the full three-year period. As noted above, estimation of the model was carried out using data sets restricted to at- and out-of-the-money options. These data sets are designated as "In-the-Money Excluded" in the table. The time-to-maturity restriction of equation (20) was tested by estimating an unrestricted version of the Hermite basis pricing model that uses a single option maturity class. This was the first option maturity class on any day with time to maturity in excess of 30 days. Separate samples were drawn of two calls and two puts at 5-minute intervals for multiple maturity classes and for a single maturity class. The average minimum option maturity for the "Multiple Maturity Classes" sample is much smaller than that for the "Single Maturity Class" sample because the former sample included serially listed options in addition to the quarterly-cycle expiration options. (The CBOE lists three near-term months and three additional months from the March quarterly cycle.) Because the serial options tend to have lower volume, they were excluded from the construction of the single maturity class sample.

For the evaluation of prediction errors conducted below, two other samples were drawn that included all option strikes, with one sample having multiple maturity classes and the other limited to a single maturity class. There were 759 trading days in the three-year period. Days were lost in each sample because of extremely low volume that resulted in fewer than 10 time-series cross sections being extracted. These days were skipped if this (arbitrary) threshold was crossed. More days were omitted for the single maturity class samples because of the smaller pool of available price quotations.

Table 1 also gives the means and standard deviations of daily averages or extreme values of sample characteristics. In the "In-the-Money Excluded, Multiple Maturity Classes" sample, calls are on average 2.7 percent out of the money, whereas puts are 3.6 percent out of the money. Both calls and puts range as far as 9 percent out of the money on average. Calls and puts varied from 48 days to maturity to 290 days, with an average of just over 100 days. There were on average four maturity classes included in each day's multiple maturity classes samples.

2.2 Estimation

We estimate the drift and variance rate parameters μ and σ in the reference measure as well as the third- and fourth-order Hermite polynomial risk prices π_3 and π_4 in the change of measure density. The empirical counterpart of equation (7) is

$$\begin{aligned} C_{it} &= \sum_{k=0}^{4} \pi_k a_k(\mu,\sigma) + u_{it}^{c} \\ P_{jt} &= \sum_{k=0}^{4} \pi_k b_k(\mu,\sigma) + u_{jt}^{p}, \end{aligned} \tag{23}$$

where C_{it} is the price of the i-th call option $(i=1,2)$ and P_{jt} is the price of the j-th put option $(j=1,2)$ at time t in the daily panel of options. The a_k and b_k coefficients above correspond to the Hermite polynomial elements in equations (12) and (15).[2]

Equation (23) was estimated in restricted and unrestricted form with regard to the treatment of time to maturity. The lower-order π_k's were restricted in estimation to be: $\pi_0 = e^{-rt}$, $\pi_1 = 0$, and $\pi_2 = 0$. The restricted version of the model used the "In the Money Excluded, Multiple Maturity Class" sample and imposed the time-to-maturity restrictions of equation (20). The unrestricted version employed the "In the Money Excluded, Single Maturity Class" sample. A standard Wald statistic evaluated using the difference between restricted and unrestricted parameter vectors is used as a formal test of the time-to-maturity restrictions.[3] The Black-Scholes restriction that $\mu = r - d$, $\pi_3 = 0$, and $\pi_4 = 0$ is also tested by a Wald statistic computed separately for the restricted and unrestricted parameter estimates.

The error terms u_{it}^{c} and u_{jt}^{p} are assumed to arise for two reasons: (1) the infinite dimension basis representation of the contingent claim price is truncated to a finite dimension and (2) the market prices of traded claims are assumed to be noisy (e.g., some quotations or

[2] These coefficients were evaluated analytically using the *Mathematica* software package.
[3] See Judge et al. (1985), 215-216.

index values may be stale, or the actual intraday interest rate or expected dividend rate may deviate from the fixed daily dividend or interest rate inputs).

Equation (23) is estimated using an iterated GMM procedure. The orthogonality conditions include four pricing error vectors, i.e., the difference between the market option price and the Hermite basis model price. Two vectors are for the sampled call options in the intraday sample, and two for the sampled put options. Vectors of the corresponding ratios of strike price to index values for each option, denoted by z^i_{tn}, where $i \in \{\text{call, put}\}$ and $n \in \{\text{sample vector 1, sample vector 2}\}$, are the instruments used in estimation, giving the following eight orthogonality conditions:

$$\mathrm{E}\left[\left(u^c_{t1} \;\; u^c_{t1}z^c_{t1} \;\; u^c_{t2} \;\; u^c_{t2}z^c_{t2} \;\; u^p_{t1} \;\; u^p_{t1}z^p_{t1} \;\; u^p_{t2} \;\; u^p_{t2}z^p_{t2}\right)'\right] = \mathbf{0} \tag{24}$$

The empirical counterpart of these orthogonality conditions is of the form

$$g(\theta) = \frac{1}{T}\sum_{t=1}^{T}\left(u^c_{t1} \;\; u^c_{t1}z^c_{t1} \;\; u^c_{t2} \;\; u^c_{t2}z^c_{t2} \;\; u^p_{t1} \;\; u^p_{t1}z^p_{t1} \;\; u^p_{t2} \;\; u^p_{t2}z^p_{t2}\right)',$$

where θ is the parameter vector and T is the number of time-series cross sections in the intraday sample. The GMM procedure minimizes the quadratic form, $g'(\theta)W(\theta)g(\theta)$, with respect to a weighting matrix $W(\theta)$.[4] The standard GMM covariance matrix is

$$\frac{1}{T}\left[D'(\theta)S^{-1}(\theta)D(\theta)\right]^{-1},$$

[4] The GMM procedure was implemented using the GAUSS Constrained Optimization procedure. The only constraint imposed in estimation was that $\sigma > 0$. Because of the fourth-order truncation of the Hermite polynomials, the estimated density is not guaranteed to integrate to unity. The market prices of risk were not constrained in estimation or normalized after estimation to enforce this restriction.

Both the unrestricted and restricted estimation runs were repeated several times for each data set using randomly perturbed starting values. Output from each run through the 3-year sample was merged by selecting the result for a given day that had the smallest *J*-statistic. (For a given convergence criterion, the GMM routine can sometimes "converge" to different points in the parameter space.)

where $D(\theta)$ is the Jacobian matrix $\partial g / \partial\theta$ and $S = \mathrm{E}[g(\theta)g'(\theta)]$.

3. Estimation Results

3.1 In-Sample Results

The output of the daily GMM estimates of the Hermite basis pricing model is summarized in Tables 2 and 3. Table 2 gives the results for the restricted model (with time-to-maturity restrictions imposed); Table 3 shows gives the corresponding results for the unrestricted model. The daily parameter estimates and test statistics are averaged over the full three-year sample period. The standard deviations of these estimates and statistics are also reported. This information is repeated for the subsample of outcomes for which the four-parameter Hermite model fails to be rejected at the 5 percent level by the GMM *J*-statistic. Eleven days from the sample were dropped in computing the results for Table 3, because the estimates for these days are clearly outliers. These days had extreme estimated skewness or excess kurtosis, which greatly distort the sample averages of the daily parameter estimates as well as the shape of the density, as depicted in Figure 1 below.[5] (The conclusions of this paper are not altered by including these dates. The single maturity class results for Table 3 include all sample dates for which estimates were obtained.)

The goodness-of-fit criterion, the *J*-statistic, fails to reject the model in 80 percent of the sample days. The daily Wald test of the Black-Scholes restrictions is usually strongly rejected. Consistent with other research on the SPX options, noted in the introduction, the risk-neutral distribution shows significant negative skewness and excess kurtosis, i.e. π_3 and π_4 are significantly negative and positive, respectively, on average. (Appendix A in Abken, Madan and

---

[5] The dates are 2/14/90, 12/17/90, 5/16/91, 8/27/91, 10/18/91, 10/22/91, 11/14/91, 12/3/91, 12/10/91, 12/17/91, and 11/9/92. The market price of skewness $\pi_3 < -5$ for all and $\pi_3 < -10$ for seven of the eleven dates. The joint impact of the estimates from these dates on the estimated risk-neutral density (graphed in Figure 1) below is most evident in the tails, which dip below zero with these dates included.

Ramamurtie (1996) shows that skewness and excess kurtosis are directly proportional to the prices of risk π_3 and π_4, respectively.)

However, there is considerable dispersion in the daily parameter and summary statistic estimates, as measured by their standard deviations in Table 2. The average value of σ is high, .23, relative to its value of .17 estimated with the Hermite basis pricing model specialized to Black-Scholes, using the same data set. There are 19 dates for which σ for the restricted model exceeds 1.0 in the sample; the restricted model frequently fit the prices for individual dates poorly. The average estimated risk-neutral drift also appears to be high at .06. The average continuously compounded risk-free rate corresponding to the option maturities in the sample is 5.5 percent, while the average dividend yield is 3.2 percent.

The residuals from equation (22), i.e., the pricing errors, were tested for any remnant of the volatility smile or skew. Linear regressions were run, separately for puts and calls, of each day's pricing errors from the restricted four-parameter model on moneyness and moneyness squared. If the four-parameter model "flattens" the volatility smile, the R^2's should be zero. The four-parameter model does not fully level the smile, although the R^2's are not large. The magnitude of the pricing errors is considered in the next section.

Table 3 gives the results for a sample consisting of the shortest maturity option in the March quarterly cycle that had time to maturity greater than 30 days. The average maturity was about 73 days. The number of two call/two put option intraday panels drops from an average of 45 in the multiple maturity classes sample to 25 in this single maturity class sample.

The Hermite basis pricing models fits the single maturity sample much better than it does the multiple maturity one. Strictly speaking, the necessity of separate estimation of the model by option maturity is evidence of model misspecification, but nevertheless the model may prove to be useful at a practical level. This point is discussed in greater detail below. Formal evidence of model misspecification is given by a Wald statistic on the time-to-maturity restrictions. That test compares the differences between the daily point estimates of the Hermite pricing model's parameters from the restricted model with those from the unrestricted model. Although the restricted (multiple maturity class sample) *J*-statistics accept the model in 80 percent of the sample days, the Wald tests, evaluated using the unrestricted estimates of the daily covariance matrices, indicate rejection of the restricted model in 79 percent of the sample days. The mean *p*-value reported in Table 3 is 10 percent. This failure of the GMM *J*-statistic to reject may be a reflection of its oft-cited low power.

As in the case of the restricted model, the daily Wald tests overwhelmingly reject the Black-Scholes restrictions. The point estimate for negative skewness is greater for the single maturity option sample, but that for excess kurtosis is somewhat lower; however, the dispersion of these daily estimates is much smaller for the single maturity sample. The volatility skew is known to attenuate at longer maturities, which may account for the shift in the skewness parameter between these samples.[6] For the same reason, the put R^2 is almost twice as large in the single maturity class sample as it is in the multiple maturity class sample.

The .19 estimate of σ is much closer to that of the Black-Scholes value of .17, derived from the Hermite model. The risk-neutral drift of .02 looks much more reasonable and is much more tightly estimated. Its standard deviation is merely .02 compared with .20 in Table 2 for the unrestricted model. Acceptance of the Hermite basis pricing model at the 5 percent significance level by the *J*-statistic rises to 95 percent of the sample days, notwithstanding the earlier comment about low power. The prediction error analysis discussed next in fact strongly supports this specification.

---

[6] As indicated in Table 2, the multiple maturity class sample has about 42 percent of its options concentrated at the shortest maturity (which is 48 days compared with 73 for the single maturity class sample).

3.2 Prediction Error Evaluation

Table 4 gives a summary of one-day ahead prediction errors by the four-parameter Hermite basis pricing model as well by the Black-Scholes version of this model. The prediction error is the pricing error resulting from evaluating a model using the previous day's parameter values. It is important to note that whereas the model was fit in sample using option samples that excluded in-the-money options, the prediction error evaluation was performed by pricing samples that included all option strikes. Black-Scholes is taken as the benchmark model, and was evaluated simply by setting $\mu = r - d$, $\pi_3 = 0$, and $\pi_4 = 0$ in the Hermite basis pricing model. To ensure comparability between models, the Black-Scholes σ was estimated using the same set of instruments as in the estimation of the four-parameter Hermite model.

It is clear from Table 4 that the restricted Hermite basis pricing model estimated on multiple maturity classes performed poorly compared with the Black-Scholes version of the model. The Hermite model appeared to yield unstable parameter estimates for a subset of the sample days, even after screening out several obvious outliers, noted earlier. These relatively infrequent large shifts in the parameters coincide with big contemporaneous as well as one-day-ahead pricing errors.[7] The largest individual pricing errors tend to occur for longer maturity options. In contrast, the Black-Scholes model is much more robust and stable in its estimates. Because the Black-Scholes version of the model is applied to pricing the same set of options with dramatically better results, the poor contemporaneous pricing and prediction error results of the restricted four-parameter model cannot be attributed to a severe stale price problem in the quotes for the longest maturity options. The key deficiency of the four-parameter Hermite model seems to be the imposition of the time-to-maturity restrictions on the market prices of risk.

---

[7] Dates that involve a switch in option maturity class to the next available class, i.e., when the current shortest maturity class drops to 29 days, were excluded from the prediction error evaluations.

Table 5 shows the corresponding results for the single maturity class. The average daily mean error over the sample shrinks to .02 point ($2 per contract) and the mean prediction error to .03, with calls being undervalued and puts overvalued by the four-parameter Hermite basis pricing model. This is a huge improvement on the restricted model results. Failing to account for skewness and excess kurtosis using the Black-Scholes version of the model results in an average daily mean error and mean prediction error approximately ten times as large as those for the unrestricted four-parameter model. The average daily mean absolute and mean square prediction errors for the unrestricted four-parameter model are about the same size as the average bid-ask spread in the sample (see Table 1).

3.3 Implied S&P 500 Index

To gauge the impact of staleness of the reported S&P 500 index on the estimation results, separate data sets were generated for both estimation and prediction that used implied index values inferred by cost of carry from S&P 500 index futures prices. As noted above, staleness is particularly likely to be an issue during the first half hour of trading; hence this time period is excluded from the sample. However, it may also be a concern during "fast" moving markets.

Our futures data set from the Chicago Mercantile Exchange (CME) covers all trading days in 1992. The CME records futures transactions when they occur at a different price from that of the last transaction. The implied index value from the reported nearby S&P 500 index future transaction was aligned with option price quotations, provided that the futures transaction occurred no earlier than five minutes before the option quotation. If such a futures price was not available in that time frame, the option quotation was not included. Measurement error is introduced by applying a cost of carry adjustment because our interest rate series is observed daily, not intraday, and the dividend yield series is derived from a quarterly average ex post time series. The selection of the nearby contract mitigates the impact of this measurement error because the discounting of the futures prices is over a relatively short time horizon.

Using the same format as that in Table 3 for the full sample, Table 6 gives the results for 1992 using an implied index value in the pricing models. Because of the rejection of the restricted model, only the unrestricted model was evaluated using this revised data set. Table 7 shows the results for the previous single maturity class data set, evaluated only for 1992. The sample characteristics as well as the results are quite similar between the tables. The parameter estimates are close, with the greatest difference being a somewhat lower market price of excess kurtosis π_4 for the implied index results (.20 for the implied index compared with .23 for the reported index). However, the standard deviations of the daily estimates is higher for the parameter estimates in Table 6, especially the standard deviation of π_4 (.59 in Table 6; .25 in Table 5). The diagnostic regressions of the pricing errors on moneyness and moneyness squared show higher R^2's for the implied index results as well.

The differences in the in-sample results for the models based on these data sets carry over to their prediction errors. Both the Hermite basis pricing model and the Black-Scholes version of the model exhibit better average daily prediction errors and lower standard deviations of those averages for the original data set with the reported index. In Table 8, based on the implied index, calls are substantially underpriced by the Hermite model, while puts are overpriced. In contrast, the results for the reported index in Table 9 indicate the same pattern, although the degree of call underpricing is smaller. These effects are consistent with the volatility skew being more pronounced in the in-sample pricing errors from the implied index diagnostic regression results.

3.4 Interpretation

Our reference measure is one under which the underlying primary asset value process is geometric Brownian motion with drift μ and variance rate σ. Initially, we assume that μ and σ are *constant* for the entire sample time horizon such that there is a well-specified relationship between the basis prices π_k for different times to maturity. This is equation (20) above, repeated here:

$$\pi_k(s) = \pi_k(t)e^{r(t-s)}(s/t)^{k/2},$$

for $s < t$ and where k is the k-th basis element.

The restrictions have been derived in a preference-free setting commonly employed in option pricing models. Equivalently, as shown in Appendix C of Abken, Madan, and Ramamurtie (1996), this relationship is consistent with an economy in which investors maximize their terminal consumption or wealth and their preferences do not depend on the price path of the underlying asset.[8]

In an asset pricing context, the investors' marginal utilities depend only on the final asset price, not the path followed by that price. In other words, empirical rejection of (20) could imply that investors' utilities may depend on intermediate consumption, even when the parameters μ and σ are constant. This implies that the relationship between basis prices at different dates is more complex than that underlying the derivation of the time-to-maturity restrictions.

An alternative reason for the rejection of the deterministic restrictions of equation (20) could be that μ and σ do vary over time as functions of some other state variables. In fact, Madan and Milne's general model does specify that μ and σ could be functions of another Markov process $\xi(t)$. With time-varying μ and σ, we would have to explicitly specify the $\xi(t)$ process and the functional dependence of the parameters on that process in order to derive restrictions similar to (20), which could then be empirically verified. Under this assumption, we make no claims about investor preferences.

In our empirical tests, we retained the constant parameter assumption and also estimated an unrestricted model, i.e., one without restriction (20). A justification for the unrestricted model is that contingent claims can be arranged in different maturity classes and that each maturity class

---

[8] This appendix gives the technical conditions on the risk-neutral and reference measures in order for equation (20) to hold.

is driven by a *separate*, but possibly correlated, Wiener process, with similar Hermite polynomial representations. As discussed above, the unrestricted model performed well.

We recognize the need for a more general setup that incorporates time-varying parameters, with corresponding relationships between the parameters for different maturity classes. The unrestricted model effectively implies separate Wiener processes for different option maturities. However, all of the options are claims on the same underlying asset; therefore, the treatment of the unrestricted model appears to be ad hoc.

However, while there may be only one underlying Wiener process driving the primary asset value *directly*, there may be several other Wiener processes (through $\xi(t)$) that affect the asset values through their impact on μ and σ. The unrestricted four-parameter Hermite model results may be giving indirect evidence on the impact of the state variable or stochastic process $\xi(t)$, with its own embedded discrete-time representation.

Thus, estimation of the Hermite model parameters separately for different maturities is defensible. The fact that the μ's and σ's, especially the latter, are different for different maturities makes the unrestricted approach more reasonable and more valuable in practical applications than the full restricted model.

3.5 Statistical Density Estimation

The statistical density given by equation (21) was estimated by maximum likelihood using daily closing values for the S&P 500 index. The 1990–1992 sample standard deviation, σ, of 16 percent reported in Table 10 exceeds slightly the average 15 percent value of the daily σ estimates from the risk-neutral density estimation. The statistical drift during this period was 11 percent; however, this was estimated imprecisely and is insignificantly different from zero. Skewness in the statistical density is also zero, while excess kurtosis is significant and somewhat greater than that estimated for the risk-neutral density in Table 3 (.23 statistical; .19 risk-neutral).

3.6 Density Comparisons

The upper panels in Figures 1 and 2 graph the estimated risk-neutral densities derived from the multiple and single maturity class data sets, respectively. Each upper panel also shows identical plots statistical density discussed in the previous section as well as the change of measure density (i.e., $\lambda(z)/\upsilon(z)$ from equation (1) above) in the lower panel. Qualitatively, the risk-neutral densities are similar in both figures.

The upper panel in each figure includes plots of risk-neutral density derived from the four-parameter Hermite basis pricing model and one for the single-parameter, Black-Scholes version of the model. Because the horizontal axis is simply the standardized normal variate, the Black-Scholes density is just the standard normal density. The other graphs were generated by evaluating the risk-neutral density (19) using the sample averages of the daily estimated π_3 and π_4 prices of risk and statistical density (21) using the maximum likelihood estimates for α_3 and α_4. The Hermite model's risk-neutral density shows excess kurtosis and a thicker left-hand tail compared with the normal. A greater degree of skewness in the risk-neutral density is evident in Figure 2's plots for the unrestricted model results, which had a shorter average maturity than that for the restricted model.

Because the Hermite model was truncated at fourth order, there is an approximation error. This is apparent in Figure 2 for the unrestricted model results. The right-hand tail of the density implied by the unrestricted model dips slightly below zero beyond two standard deviations from the mean.[9] Both the symmetry of the statistical density as well as its excess kurtosis are readily seen in comparison with the standard normal density in each figure.

The change of measure density in each Figure gives a clearer view of how the symmetric statistical density is altered to give the left-skewed risk-neutral density. This is a graphical illustration of the relative premium for out-of-the-money S&P 500 index put options vis-à-vis out-of-the-money call options.

4. Conclusion

This paper tests the approach of Madan and Milne (1994) and Abken, Madan, and Ramamurtie (1996) for pricing contingent claims as elements of a separable Hilbert space. A key component of the model is the restrictions on the prices of Hermite polynomial risks for contingent claims with different times to maturity. These restrictions, which constrain the shape of the risk-neutral density at future points in time, allow all option maturity classes to be used in estimation of the model. The restrictions are rejected by our empirical tests of the four-parameter model on the S&P 500 index options.

We drop the market price of risk restrictions and estimate the model separately for a single option maturity class, the most liquid in our sample. This can be justified as an indirect way to accommodate the likely time variation of the underlying parameters of the asset price dynamics, which the current model does not explicitly include. The out-of-sample performance of the unrestricted four-parameter model is consistently better than that of the unrestricted Black-Scholes version of the model.

The price of risk estimates for the restricted as well as unrestricted models indicate skewness and excess kurtosis in the implied risk-neutral density. On the other hand, the estimated statistical density is symmetric. It displays excess kurtosis of about the same magnitude as the risk-neutral density's. This contrasts with our previous empirical work that applied the Hermite basis pricing model to valuing Eurodollar futures options. Skewness of the implied risk-neutral distribution for the SPX options is much more pronounced and consistently in one direction.

---

[9] The prediction error analysis reported in Table 5 indicates that calls are on average underpriced by the model. The negative area in the right-hand tail may contribute to this underpricing. However, note from Table 1 that the average call in the prediction sample "All Strikes, Single Maturity Class" is almost 3 percent in the money.

Another contrast is that, unlike the implicit distributions for the SPX options and S&P 500 index, the risk-neutral and statistical densities for short-term Eurodollar options are quite similar in shape, both exhibiting statistically significant skewness and excess kurtosis.

Future work can test the static hedging properties of the model. In particular, the Hermite polynomial coefficients have the interpretation of being basis-mimicking portfolios. In addition, by deriving the Hermite polynomial coefficients for a set of contingent claims and estimating their market prices of risk, arbitrage-free pricing of other, possibly less liquid contingent claims is feasible by Hermite polynomial approximation. Many contracts, such as path-dependent options, exotic options, swaps, etc., are candidates for this approach.

References

Abken, Peter A., Dilip B. Madan, and Sailesh Ramamurtie. "Estimation of Risk-Neutral and Statistical Densities By Hermite Polynomial Approximation: With an Application to Eurodollar Futures Options." Federal Reserve Bank of Atlanta Working Paper 96-5, June 1996.

Aït-Sahalia, Yacine, and Andrew W. Lo. "Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices." NBER Working Paper 5351, November 1995.

Bates, David S. "The Crash of '87: Was It Expected? The Evidence from Options Markets." *Journal of Finance* 46 (July 1991): 1009-44.

Bossaerts, Peter, and Pierre Hillion. "A Test of a General Equilibrium Stock Option Pricing Model." *Mathematical Finance* 3 (October 1993): 311-347.

Cox, John C., and Mark Rubinstein. *Option Markets*. Englewood Cliffs: Prentice-Hall, Inc., 1985.

Derman, Emanuel, and Iraj Kani. "Riding on a Smile." *Risk* 7 (February 1994): 32-39.

Duffie, Gregory R. "Idiosyncratic Variation in Treasury Bill Yields." *Journal of Finance* 51 (June 1996): 527-552.

Dumas, Bernard, Jeff Fleming, and Robert E. Whaley. "Implied Volatility Functions: Empirical Tests." NBER Working Paper #5500, March 1996.

Dupire, Bruno. "Pricing with a Smile." *Risk* 7 (January 1994): 18-20.

Dwyer, Gerald P., Jr., Peter Locke, and Wei Yu. "Index Arbitrage and Nonlinear Dynamics Between the S&P 500 Futures and Cash." *Review of Financial Studies* 9 (1996): 301-332.

Elliot, Robert J. "A General Recursive Discrete Time Filter." *Journal of Applied Probability* 30 (1993): 575-588.

Jackwerth, Jens Carsten, and Mark Rubinstein. "Implied Probability Distributions: Empirical Analysis." Institute of Business and Economic Research Finance Working Paper #250, University of California at Berkeley, June 1995.

Judge, George G., W.E. Griffiths, R. Carter Hill, Helmut Lütkepohl, and Tsoung-Chao Lee. *The Theory and Practice of Econometrics.* 2nd edition. New York: John Wiley and Sons, 1985.

Longstaff, Francis A. "An Empirical Examination of the Risk-Neutral Valuation Model." Working Paper, College of Business, Ohio State University, and the Anderson Graduate School of Management, UCLA, 1992.

__________. "Option Pricing and the Martingale Restriction." *Review of Financial Studies* 8 (Winter 1995): 1091-1124.

Madan, Dilip B., and Frank Milne. "Contingent Claims Valued and Hedged by Pricing and Investing in a Basis." *Mathematical Finance* 4 (July 1994): 223-245.

Rozanov, Y. A. *Markov Random Fields*. New York: Springer Verlag, 1982.

Rubinstein, Mark. "Implied Binomial Trees." *Journal of Finance* 49 (July 1994): 771-818.

Shimko, David. "Bounds of Probability." *Risk* 6 (April 1993): 33-37.

Table 1
Descriptive Statistics for the Sampled S&P 500 Index Options January 2, 1990 — December 31, 1992

	In-the-Money Excluded Multiple Maturity Classes		*In-the-Money Excluded Single Maturity Class*		*All Strikes Multiple Maturity Classes*		*All Strikes Single Maturity Class*	
Days in sample	749		643		755		737	
	Mean	*Std. Dev.*	*Mean*	*Std. Dev.*	*Mean*	*Std. Dev.*	*Mean*	*Std. Dev.*
Average option price	8.68	1.85	7.11	2.45	15.52	2.64	16.17	3.32
Average bid-ask spread	0.43	0.08	0.37	0.09	0.54	0.09	0.54	0.10
Average index level	374.89	36.49	371.59	35.64	375.28	36.61	374.14	36.26
Average call moneyness	2.73	0.75	2.68	0.82	-1.79	1.70	-2.60	1.92
Maximum call moneyness	8.96	1.72	8.18	1.95	8.77	1.80	7.92	2.30
Minimum call moneyness	0.18	0.24	0.23	0.26	-8.61	0.74	-8.62	0.68
Average put moneyness	-3.59	0.65	-3.47	0.85	-0.51	1.25	-0.35	1.59
Maximum put moneyness	-0.22	0.26	-0.31	0.39	8.45	1.99	7.20	2.62
Minimum put moneyness	-8.57	0.58	-8.34	0.81	-8.48	0.67	-8.27	0.87
Average call maturity	107.60	23.40	72.72	26.57	89.94	22.30	75.23	26.45
Maximum call maturity	285.17	47.58	72.72	26.57	283.79	49.47	75.23	26.45
Minimum call maturity	47.88	14.33	72.72	26.57	48.03	14.38	75.23	26.45
Average put maturity	104.36	23.13	72.72	26.57	91.11	20.98	75.23	26.45
Maximum put maturity	290.83	44.14	72.72	26.57	290.13	45.50	75.23	26.45
Minimum put maturity	47.88	14.33	72.72	26.57	48.03	14.38	75.23	26.45
Av. No. of call expirations	4.07	0.65	1	0	4.05	0.66	1	0
Av. No. of put expirations	4.14	0.66	1	0	4.11	0.67	1	0

Average and extreme values are computed for each day's sampled options. The means and standard deviations of these daily values are then computed over the three-year sample. The first two out-of-the-money calls and two out-of-the-money puts in each five-minute interval during the trading day are sampled. Moneyness is $\left[(\text{strike} / \text{index}) - 1\right] \times 100$ and maturity is expressed in days. The option price is expressed in points of 100; the actual price of a contract is this value times \$100. "Multiple Maturity Classes" samples include serially listed expiration months, whereas "Single Maturity Class" samples are restricted to the more liquid quarterly cycle expiration months. The "In-the-Money Excluded" samples are used for estimation of the Hermite basis pricing model. The "All Strikes" samples are used for evaluation of the prediction errors of the model.

Table 2

In-Sample Estimates for the Hermite Basis Pricing Model
S&P500 Index Options
Multiple Maturity Classes

January 2, 1990 — December 31, 1992

	Full Sample		*J-statistic p-value > .05*	
Sample Size	738		589	
	Mean	*Std. Dev.*	*Mean*	*Std. Dev.*
J-statistic	6.80	4.57	4.97	2.26
J-statistic p-value	0.29	0.26	0.36	0.25
Wald statistic p-value for Black-Scholes restrictions	0.003	0.03	0.004	0.03
μ	0.06	0.20	0.06	0.22
σ	0.23	0.22	0.24	0.23
π_3	-0.30	0.49	-0.27	0.43
π_4	0.23	0.87	0.19	0.81
Average call moneyness	2.72	0.75	2.68	0.73
Maximum call moneyness	8.96	1.72	8.94	1.78
Min call moneyness	0.18	0.24	0.18	0.24
Average put moneyness	-3.59	0.65	-3.61	0.66
Maximum put moneyness	-0.21	0.26	-0.21	0.26
Minimum put moneyness	-8.58	0.58	-8.60	0.58
Average call maturity	107.58	23.31	105.38	22.27
Maximum call maturity	285.16	47.70	285.97	49.19
Minimum call maturity	47.85	14.11	47.75	13.45
Average put maturity	104.64	23.06	104.47	23.10
Maximum put maturity	291.39	43.77	291.61	44.69
Minimum put maturity	47.85	14.11	47.75	13.45
Intraday sample size	45.60	12.47	44.76	12.58
Proportion of shortest maturity	41.86	14.11	42.36	14.28
Call R^2	0.19	0.21	0.17	0.19
Put R^2	0.09	0.10	0.08	0.09

The shortest maturity option on a given day was restricted to have time-to-maturity in excess of 30 days. This sample included serial and quarterly expiration maturity classes. Average and extreme values are computed for each day's sampled options. The means and standard deviations of these daily values are then computed over the three-year sample. Moneyness is $[(\text{strike}/\text{index})-1]\times 100$ and maturity is expressed in days. The intraday sample size is the number of 5-minute intervals included that two calls and two puts meeting the selection criteria. The call and put R^2 entries give the R^2 values for regressions of the pricing errors on a constant and option moneyness and moneyness squared. The Wald statistic tests the restriction that $\mu = r - d$, $\pi_3 = 0$, and $\pi_4 = 0$.

Table 3

In-Sample Estimates for the Hermite Basis Pricing Model
S&P500 Index Options
Single Maturity Class

January 2, 1990 — December 31, 1992

	Full Sample		*J-statistic p-value > .05*	
Sample Size	643		610	
	Mean	*Std. Dev.*	*Mean*	*Std. Dev.*
J-statistic	4.50	2.94	4.06	2.09
J-statistic p-value	0.43	0.27	0.46	0.25
Wald statistic p-value for Black-Scholes restrictions	< .001	< .001	< .001	< .001
Wald statistic p-value for time-to-maturity restrictions	0.10	0.24	—	—
μ	0.02	0.02	0.02	0.02
σ	0.19	0.05	0.19	0.05
π_3	-0.52	0.18	-0.52	0.15
π_4	0.19	0.32	0.20	0.30
Average call moneyness	2.68	0.82	2.68	0.82
Maximum call moneyness	8.18	1.95	8.14	1.96
Minimum call moneyness	0.23	0.26	0.23	0.26
Average put moneyness	-3.47	0.85	-3.48	0.85
Maximum put moneyness	-0.31	0.39	-0.31	0.39
Minimum put moneyness	-8.34	0.81	-8.34	0.81
Average call maturity	72.72	26.57	72.66	26.44
Maximum call maturity	72.72	26.57	72.66	26.44
Minimum call maturity	72.72	26.57	72.66	26.44
Average put maturity	72.72	26.57	72.66	26.44
Maximum put maturity	72.72	26.57	72.66	26.44
Minimum put maturity	72.72	26.57	72.66	26.44
Intraday sample size	25.76	13.32	25.43	13.31
Proportion of shortest maturity	100	0	100	0
Call R^2	0.18	0.22	0.17	0.20
Put R^2	0.15	0.23	0.13	0.20

The option maturity class is that of the shortest maturity option (in the quarterly expiration cycle) on a given day with time-to-maturity in excess of 30 days. Average and extreme values are computed for each day's sampled options. The means and standard deviations of these daily values are then computed over the three-year sample. Moneyness is $[(\text{strike} / \text{index}) - 1] \times 100$ and maturity is expressed in days. The intraday sample size is the number of 5-minute intervals included that two calls and two puts meeting the selection criteria. The call and put R^2 entries give the R^2 values for regressions of the pricing errors on a constant, option moneyness, and moneyness squared. The Black-Scholes Wald statistic tests the restriction that $\mu = r - d$, $\pi_3 = 0$, and $\pi_4 = 0$, while the Wald statistic for the time-to-maturity restrictions tests for the equality of the estimated parameter vectors of the restricted and unrestricted Hermite models.

Table 4

Evaluation of Prediction Errors
Multiple Maturity Classes

January 2, 1990 — December 31, 1992

	All		*Calls*		*Puts*	
Hermite Basis Pricing Model						
Sample size: 672						
Mean error	-1.73	(45.21)	-4.15	(88.47)	0.68	(5.71)
Mean absolute error	9.32	(55.62)	13.15	(104.56)	5.50	(10.99)
Mean square error	24.38	(196.85)	30.45	(277.72)	10.62	(22.78)
Mean prediction error	-2.22	(41.35)	-4.69	(77.51)	0.25	(7.85)
Mean prediction error Outside bid-ask spread	-2.19	(41.35)	-4.67	(77.51)	0.28	(7.81)
Mean absolute prediction error	10.09	(54.59)	14.39	(99.63)	5.78	(12.51)
Mean square prediction error	26.15	(175.47)	33.14	(247.05)	10.87	(26.39)
Black-Scholes Model						
Sample size: 674						
	All		*Calls*		*Puts*	
Mean error	0.35	(0.45)	0.48	(0.50)	0.22	(0.56)
Mean absolute error	1.13	(0.28)	1.18	(0.32)	1.08	(0.32)
Mean square error	1.41	(0.33)	1.44	(0.35)	1.36	(0.38)
Mean prediction error	0.35	(0.53)	0.48	(0.57)	0.22	(0.63)
Mean prediction error Outside bid-ask spread	0.28	(0.47)	0.32	(0.49)	0.24	(0.55)
Mean absolute prediction error	1.16	(0.30)	1.21	(0.34)	1.10	(0.33)
Mean square prediction error	1.44	(0.36)	1.47	(0.38)	1.39	(0.41)

The table gives the means and standard deviations of the daily error measurements as defined in the left-hand column. The sample of options being priced includes all option strikes. The mean error, mean absolute error, and mean square error for a given day include in-the-money options that were not used in estimating model parameters. The prediction error is the pricing error resulting from evaluating a model using the previous trading day's parameter values. The computation of prediction errors was restricted to days for which parameter estimates were available for the previous trading day. The mean prediction error outside the bid-ask spread sets pricing errors that fall within the bid-ask spread to zero (see Dumas, Fleming, and Whaley 1996).

Table 5

Evaluation of Prediction Errors
Single Maturity Class

January 2, 1990 — December 31, 1992

	All		*Calls*		*Puts*	
	Hermite Basis Pricing Model					
	Sample size: 580					
Mean error	0.02	(0.19)	0.08	(0.45)	-0.03	(0.24)
Mean absolute error	0.31	(0.30)	0.36	(0.40)	0.27	(0.24)
Mean square error	0.44	(0.59)	0.46	(0.52)	0.39	(0.68)
Mean prediction error	0.03	(0.40)	0.09	(0.56)	-0.02	(0.51)
Mean prediction error Outside bid-ask spread	0.03	(0.31)	0.06	(0.44)	0.00	(0.40)
Mean absolute prediction error	0.44	(0.38)	0.46	(0.44)	0.42	(0.39)
Mean square prediction error	0.57	(0.63)	0.55	(0.49)	0.54	(0.77)
	Black-Scholes Model					
	Sample size: 580					
	All		*Calls*		*Puts*	
Mean error	0.26	(0.39)	0.48	(0.56)	0.04	(0.47)
Mean absolute error	1.01	(0.34)	1.15	(0.45)	0.87	(0.31)
Mean square error	1.19	(0.57)	1.30	(0.49)	1.05	(0.68)
Mean prediction error	0.25	(0.53)	0.46	(0.64)	0.04	(0.62)
Mean prediction error Outside bid-ask spread	0.18	(0.46)	0.28	(0.55)	0.09	(0.53)
Mean absolute prediction error	1.04	(0.38)	1.16	(0.47)	0.91	(0.36)
Mean square prediction error	1.22	(0.59)	1.30	(0.48)	1.09	(0.74)

The table gives the means and standard deviations of the daily error measurements as defined in the left-hand column. The sample of options being priced includes all option strikes. The mean error, mean absolute error, and mean square error for a given day include in-the-money options that were not used in estimating model parameters. The prediction error is the pricing error resulting from evaluating a model using the previous trading day's parameter values. The computation of prediction errors was restricted to days for which parameter estimates were available for the previous trading day. The mean prediction error outside the bid-ask spread sets pricing errors that fall within the bid-ask spread to zero (see Dumas, Fleming, and Whaley 1996).

Table 6

In-Sample Estimates for the Hermite Basis Pricing Model Using S&P500 Index Futures to Derive Implied Index Single Maturity Class

January 2, 1992 — December 31, 1992

	Full Sample		*J-statistic p-value > .05*	
Sample Size	181		171	
	Mean	*Std. Dev.*	*Mean*	*Std. Dev.*
J-statistic	4.79	3.61	4.22	2.08
J-statistic p-value	0.41	0.26	0.44	0.25
Wald statistic p-value for Black-Scholes restrictions	< .001	< .001	< .001	< .001
μ	0.001	0.006	0.002	0.006
σ	0.16	0.03	0.16	0.03
π_3	-0.45	0.15	-0.45	0.14
π_4	0.20	0.59	0.20	0.58
Average call moneyness	2.40	0.67	2.40	0.68
Maximum call moneyness	7.58	1.77	7.54	1.75
Minimum call moneyness	0.17	0.27	0.17	0.28
Average put moneyness	-3.48	0.75	-3.49	0.75
Maximum put moneyness	-0.29	0.35	-0.30	0.34
Minimum put moneyness	-8.40	0.86	-8.39	0.87
Average call maturity	68.06	25.33	67.52	25.08
Maximum call maturity	68.06	25.33	67.52	25.08
Minimum call maturity	68.06	25.33	67.52	25.08
Average put maturity	68.06	25.33	67.52	25.08
Maximum put maturity	68.06	25.33	67.52	25.08
Minimum put maturity	68.06	25.33	67.52	25.08
Intraday sample size	28.55	18.63	27.87	18.24
Proportion of shortest maturity	100	0	100	0
Call R^2	0.23	0.26	0.22	0.26
Put R^2	0.26	0.33	0.24	0.32

The index level was derived from S&P500 index futures contracts. The option expiration is that of the shortest maturity option (in the quarterly expiration cycle) on a given day with time-to-maturity in excess of 30 days. Average and extreme values are computed for each day's sampled options. The means and standard deviations of these daily values are then computed over the three-year sample. Moneyness is $[(\text{strike / index}) - 1] \times 100$ and maturity is expressed in days. The intraday sample size is the number of 5-minute intervals included that two calls and two puts meeting the selection criteria. The call and put R^2 entries give the R^2 values for regressions of the pricing errors on a constant, option moneyness, and moneyness squared. The Wald statistic tests the restriction that $\mu = r - d$, $\pi_3 = 0$, and $\pi_4 = 0$.

Table 7

In-Sample Estimates for the Hermite Basis Pricing Model Using S&P500 Index Single Maturity Class

January 2, 1992 — December 31, 1992

	Full Sample		*J-statistic p-value* > .05	
Sample Size	181		166	
	Mean	*Std. Dev.*	*Mean*	*Std. Dev.*
J-statistic	4.67	3.75	3.89	2.12
J-statistic p-value	0.44	0.28	0.48	0.27
Wald statistic p-value for Black-Scholes restrictions	< .001	< .001	< .001	< .001
μ	0.003	0.008	0.003	0.008
σ	0.15	0.02	0.15	0.02
π_3	-0.47	0.13	-0.47	0.11
π_4	0.23	0.25	0.25	0.24
Average call moneyness	2.43	0.66	2.42	0.67
Maximum call moneyness	7.59	1.75	7.54	1.74
Minimum call moneyness	0.22	0.25	0.22	0.25
Average put moneyness	-3.46	0.75	-3.46	0.74
Maximum put moneyness	-0.31	0.32	-0.30	0.33
Minimum put moneyness	-8.36	0.84	-8.36	0.87
Average call maturity	68.22	25.31	68.55	25.18
Maximum call maturity	68.22	25.31	68.55	25.18
Minimum call maturity	68.22	25.31	68.55	25.18
Average put maturity	68.22	25.31	68.55	25.18
Maximum put maturity	68.22	25.31	68.55	25.18
Minimum put maturity	68.22	25.31	68.55	25.18
Intraday sample size	28.61	18.71	28.07	18.87
Proportion of shortest maturity	100	0	100	0
Call R^2	0.17	0.23	0.16	0.22
Put R^2	0.17	0.25	0.15	0.22

The option expiration is that of the shortest maturity option (in the quarterly expiration cycle) on a given day with time-to-maturity in excess of 30 days. Average and extreme values are computed for each day's sampled options. The means and standard deviations of these daily values are then computed over the three-year sample. Moneyness is $[(\text{strike} / \text{index}) - 1] \times 100$ and maturity is expressed in days. The intraday sample size is the number of 5-minute intervals included that two calls and two puts meeting the selection criteria. The call and put R^2 entries give the R^2 values for regressions of the pricing errors on a constant, option moneyness, and moneyness squared. The Wald statistic tests the restriction that $\mu = r - d$, $\pi_3 = 0$, and $\pi_4 = 0$.

Table 8

Evaluation of Prediction Errors
Using S&P500 Index Futures to Derive Implied Index .
Single Maturity Class

January 2, 1992 — December 31, 1992

	Hermite Basis Pricing Model					
	Sample size: 152					
	All		*Calls*		*Puts*	
Mean error	0.08	(0.30)	0.19	(0.55)	-0.03	(0.18)
Mean absolute error	0.30	(0.39)	0.34	(0.52)	0.26	(0.29)
Mean square error	0.40	(0.51)	0.45	(0.64)	0.33	(0.35)
Mean prediction error	0.04	(0.34)	0.11	(0.49)	-0.03	(0.31)
Mean prediction error Outside bid-ask spread	0.03	(0.27)	0.06	(0.40)	-0.01	(0.21)
Mean absolute prediction error	0.33	(0.33)	0.36	(0.42)	0.30	(0.25)
Mean square prediction error	0.41	(0.43)	0.44	(0.52)	0.37	(0.33)
	Black-Scholes Model					
	Sample size: 152					
	All		*Calls*		*Puts*	
Mean error	0.25	(0.26)	0.35	(0.36)	0.15	(0.37)
Mean absolute error	0.82	(0.22)	0.88	(0.29)	0.76	(0.20)
Mean square error	0.96	(0.23)	1.00	(0.29)	0.90	(0.21)
Mean prediction error	0.23	(0.39)	0.32	(0.48)	0.14	(0.45)
Mean prediction error Outside bid-ask spread	0.19	(0.32)	0.19	(0.39)	0.19	(0.37)
Mean absolute prediction error	0.83	(0.24)	0.88	(0.30)	0.78	(0.23)
Mean square prediction error	0.96	(0.27)	1.00	(0.32)	0.91	(0.25)

The table gives the means and standard deviations of the daily error measurements as defined in the left-hand column. The sample of options being priced includes all option strikes. The mean error, mean absolute error, and mean square error for a given day include in-the-money options that were not used in estimating model parameters. The prediction error is the pricing error resulting from evaluating a model using the previous trading day's parameter values. The computation of prediction errors was restricted to days for which parameter estimates were available for the previous trading day. The mean prediction error outside the bid-ask spread sets pricing errors that fall within the bid-ask spread to zero (see Dumas, Fleming, and Whaley 1996).

Table 9

Evaluation of Prediction Errors of the Basis Pricing Model
Using S&P500 Index
Single Maturity Class

January 2, 1992 — December 31, 1992

	Hermite Basis Pricing Model					
	Sample size: 153					
	All		*Calls*		*Puts*	
Mean error	0.04	(0.13)	0.10	(0.29)	-0.02	(0.19)
Mean absolute error	0.27	(0.20)	0.30	(0.24)	0.24	(0.19)
Mean square error	0.36	(0.27)	0.39	(0.31)	0.32	(0.24)
Mean prediction error	0.03	(0.28)	0.06	(0.36)	-0.002	(0.40)
Mean prediction error Outside bid-ask spread	0.02	(0.20)	0.03	(0.24)	0.02	(0.30)
Mean absolute prediction error	0.36	(0.26)	0.37	(0.25)	0.35	(0.31)
Mean square prediction error	0.46	(0.32)	0.46	(0.32)	0.43	(0.35)
	Black-Scholes Model					
	Sample size: 153					
	All		*Calls*		*Puts*	
Mean error	0.21	(0.36)	0.33	(0.47)	0.08	(0.48)
Mean absolute error	0.86	(0.28)	0.91	(0.33)	0.81	(0.28)
Mean square error	1.00	(0.31)	1.05	(0.35)	0.94	(0.30)
Mean prediction error	0.19	(0.46)	0.30	(0.58)	0.07	(0.54)
Mean prediction error Outside bid-ask spread	0.14	(0.40)	0.17	(0.50)	0.12	(0.46)
Mean absolute prediction error	0.87	(0.32)	0.92	(0.39)	0.82	(0.29)
Mean square prediction error	1.01	(0.35)	1.05	(0.41)	0.95	(0.32)

The table gives the means and standard deviations of the daily error measurements as defined in the left-hand column. The sample of options being priced includes all option strikes. The mean error, mean absolute error, and mean square error for a given day include in-the-money options that were not used in estimating model parameters. The prediction error is the pricing error resulting from evaluating a model using the previous trading day's parameter values. The computation of prediction errors was restricted to days for which parameter estimates were available for the previous trading day. The mean prediction error outside the bid-ask spread sets pricing errors that fall within the bid-ask spread to zero (see Dumas, Fleming, and Whaley 1996).

Table 10
Estimates of the Statistical Density Parameters
S&P500 Index

January 2, 1990 — December 31, 1992

Sample Size: 759

Parameter	*Estimate*	*Standard Error*
μ	0.11	0.12
σ	0.16	0.004
α_3	-0.0045	0.05
α_4	0.23	0.05

Estimation by maximum likelihood of equation (17).

Figure 1

Risk-Neutral and Statistical Densities: 1990 - 1992

Derived from S&P 500 Index and Index Options

Multiple Maturity Classes

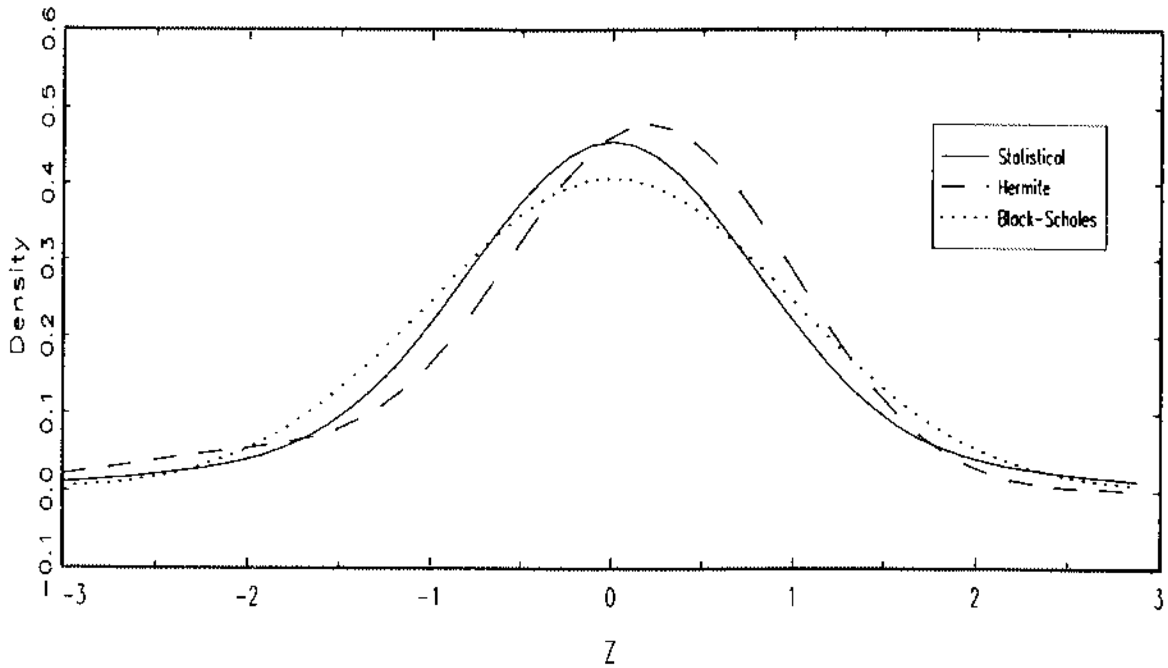

Change of Measure Density: 1990 - 1992

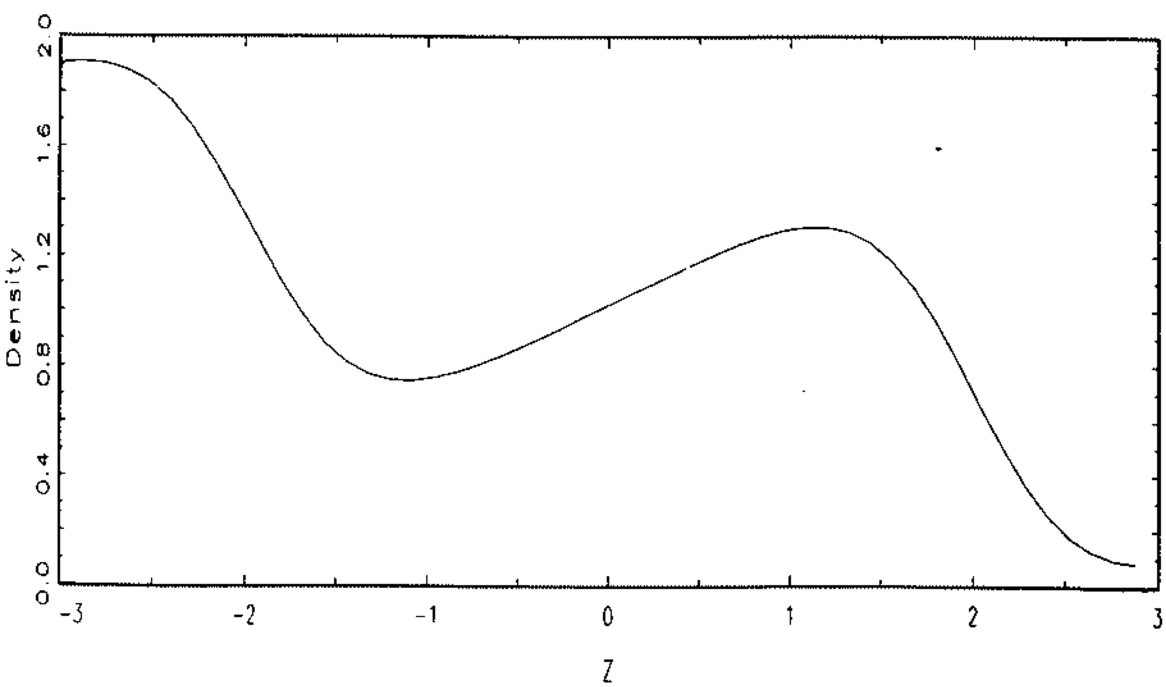

Figure 2

Risk-Neutral and Statistical Densities: 1990 - 1992

Derived from S&P 500 Index and Index Options

Single Maturity Class

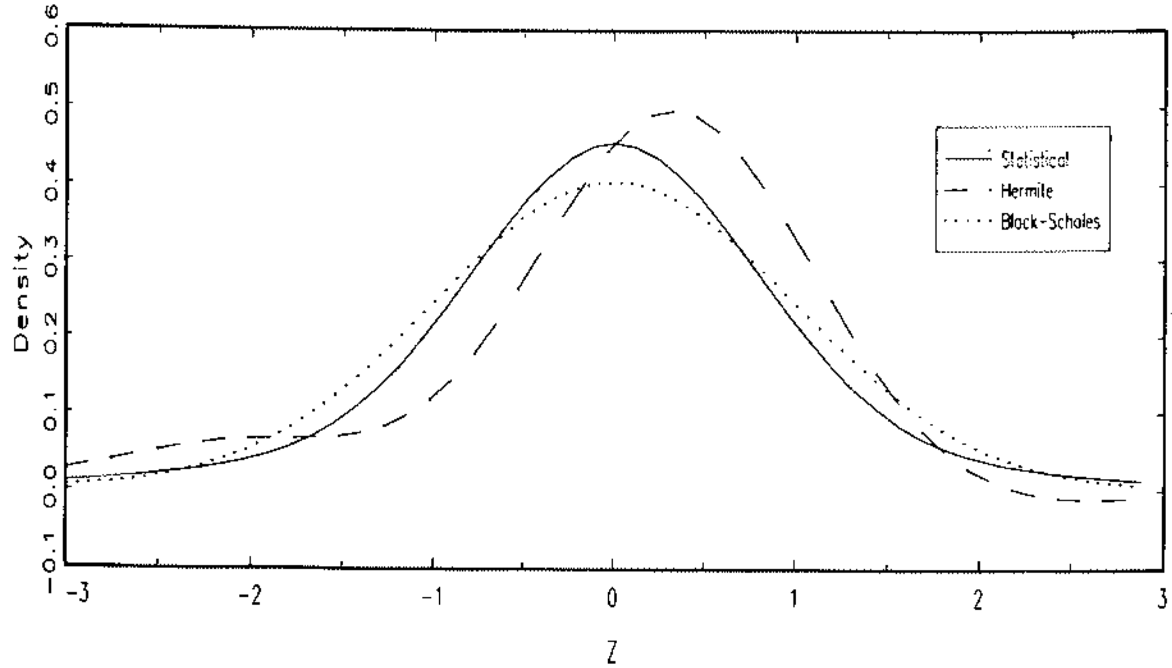

Change of Measure Density: 1990 - 1992