OPTION HEDGING USING EMPIRICAL PRICING KERNELS

Joshua V. Rosenberg
Robert F. Engle

Working Paper **6222**

NBER WORKING PAPER SERIES

OPTION HEDGING USING EMPIRICAL PRICING KERNELS

Joshua V. Rosenberg
Robert F. Engle

Working Paper 6222
http://www.nber.org/papers/w6222

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 1997

This paper is part of NBER's research program in Asset Pricing. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

Option Hedging Using Empirical Pricing Kernels
Joshua V. Rosenberg and Robert F. Engle
NBER Working Paper No. 6222
October 1997
Asset Pricing

ABSTRACT

This paper develops a method for option hedging which is consistent with time-varying preferences and probabilities. The preferences are expressed in the form of an empirical pricing kernel (EPK), which measures the state price per unit probability, while probabilities are derived from an estimated stochastic volatility model of the form GARCH components with leverage. State prices are estimated using the flexible risk-neutral density method of Rosenberg (1995) and a daily cross-section of option premia. Time-varying preferences over states are linked to a dynamic model of the underlying price to obtain a one-day ahead forecast of derivative price distributions and minimum variance hedge ratios.

Empirical results suggest that risk aversion over S&P500 return states is substantially higher than risk aversion implied by Black-Scholes state prices and probabilities using long run estimates of S&P500 return moments. It is also found that the daily level of risk aversion is strongly positively autocorrelated, negatively correlated with S&P500 price changes, and positively correlated with the spread between implied and objective volatilities.

Hedging results reveal that typical hedging techniques for out-of-the-money S&P500 index options, such as Black-Scholes or historical minimum variance hedging, are inferior to the EPK hedging method. Thus, time-varying preferences and probabilities appear to be an important factor in the day-to-day pricing of S&P500 options.

Joshua V. Rosenberg
Department of Finance
Stern School of Business
New York University
44 West 4th Street, Suite 9-190
New York, NY 10012-1126

Robert F. Engle
Department of Economics
University of California, San Diego
La Jolla, CA 92093-0508
and NBER
rengle@ucsd.edu

I. Introduction

Following the lead of Black and Scholes (1973), the three elements of asset pricing — payoffs, probabilities, and preferences — are approached in a similar manner in many option pricing models. Payoffs are determined by the contractual specifications of the asset, while probabilities are defined by a dynamic model of the underlying price. Subsequently, the underlying price process parameters are estimated from historical price data, and hedge ratios are obtained by differentiating the option pricing formula with respect to state variables. Usually, option pricing models set in discrete-time (e.g. Rubinstein (1976), Stapleton and Subrahmanyam (1982), Amin and Ng (1993), or Duan (1995)) require specific preference assumptions, but do not require estimation of preference parameters. Little attention has been focused on the realism of the implicit or explicit preference assumptions inherent in the specification of the underlying price process or solution of the model, and the effects of preference misspecification on hedge performance.

Modern approaches to option pricing including Derman and Kani (1994) and Rubinstein (1994) avoid this issue by simultaneously estimating probabilities and preferences via state prices implied by existing option premia. In these models, state prices are taken to be a deterministic function of the underlying price level and time. While these approaches are useful for interpolation of contemporaneous prices, the non-stationarity of the state-prices generated by non-stationarity of payoff probabilities suggests possible pricing instability. Dumas, Whaley, and Fleming (1996) provide evidence along these lines. Since these pricing formulas often lack a realistic model of underlying price dynamics, the hedge ratios may be inaccurate.

This paper takes a different approach to the issue of option hedging by estimating both time-varying preferences and probabilities. Probabilities are estimated using a statistical model of the underlying price process based on historical price data, and preferences are estimated using an empirical pricing kernel (EPK) which measures the state price per unit probability. The EPK is estimated as the ratio of empirical state prices and empirical state probabilities, where state prices are estimated directly from option premia. This measure of investor preferences over states is linked to a dynamic model of the underlying price to obtain hedge ratios.

A key assumption used in this paper is that for the purposes of hedging, today's pricing kernel is a sufficiently accurate estimate of tomorrow's pricing kernel. This might be considered to be a type of martingale assumption for the pricing kernel. An additional implication is that no additional risk is added to the hedge portfolio by preference changes over one day. Daily time-variation in the pricing kernel is documented in this paper, and a companion paper will examine the effects of preference variability on hedging performance.

This paper offers four primary contributions. First, a technique is developed to estimate the pricing kernel which does not require a parametric specification of the representative investor's utility function or data on aggregate consumption. The EPK technique also uses a broad set of asset prices for estimation and is updated on a daily basis using current information. Thus, it produces a daily conditional pricing kernel, rather than a long-run average pricing kernel. Non-parametric estimation of the EPK places no restrictions on the pricing kernel functional form.

The only similar work that uses option premia to estimate state prices and the pricing kernel is a paper by Jackwerth (1996). In Jackwerth's paper, the state probabilities are held constant which generates convex utility over some periods. This is likely due to impounding probabilistic information into preferences.

There is an extensive literature related to estimating pricing kernels via specification of a functional form for the utility function and estimation using aggregate consumption data. See, for example, Hansen and Singleton (1982) for a CRRA model or Cochrane and Hansen (1992) for a description of alternative utility specifications. All of these models rely on time-invariant preferences estimated using aggregate consumption data. The aggregate consumption data is reported relatively infrequently and may contain considerable measurement error. The use of aggregate consumption in the MRS also typically requires the assumption of a representative investor. In addition, identifying assumptions are usually based only on unconditional moments of stock and bond returns. An exception to this is Gallant, Hansen, and Tauchen (1990) which considers conditional moments of stock and bond returns.

A second contribution of the paper is application of a minimum variance hedging technique to the option hedging problem. Typically, option hedges are accomplished using derivatives of the pricing formula with respect to the state variables to reflect price sensitivities. This methodology is consistent with continuous hedging. In practice, hedging occurs at discrete intervals, and as Robins and Schachter (1994) have noted, the instantaneous hedge parameters will not necessarily be variance minimizing. This paper uses a simulation technique to estimate the one-day variance minimizing hedge consistent with a stochastic volatility specification for the underlying price process.

A third contribution of the paper is an empirical analysis of the performance of several methods of hedging out-of-the-money (OTM) S&P500 index put option positions. This particular application is informative since writing OTM S&P500 put options has been a historically profitable strategy, but methods of effectively hedging this position are not well known. Several hedging methods are compared including EPK hedging, Black-Scholes hedging, and a hedging technique consistent with a local volatility function.

The final contribution of the paper is an analysis of the time-series behavior of the EPKs and the implications for the historical behavior of investor preferences. This provides insight into sources of priced risk in the U.S. economy.

The paper is organized as follows. Section II describes the theory of EPK estimation and EPK hedging. Section III presents the EPK estimation technique for states defined by S&P500 returns, while section IV analyzes the hedging test results. Section V contains the conclusions.

II.a. Empirical pricing kernel theory

The purpose of this paper is to develop an option hedging methodology which combines a realistic model of asset dynamic behavior with current information about investor preferences. To this end, the EPK hedging technique utilizes a statistical model of underlying asset price behavior combined with state prices implied by option premia. The state prices reflect both investor preferences and probability forecasts, and preferences may be isolated by taking the ratio of state prices and state probabilities.

Consider the following general asset pricing equation which is consistent with an investor's solution to an optimal consumption problem. See, for example, Constantinides (1989).

(1) $$S_t = E_t\left[K_{t,T}(s)S_T(s)\right]$$

In this case, S_t is the current underlying price, $K_{t,T}(s)$ is the state price per unit probability (or pricing kernel) at date t for a payoff at date T in state s, and $S_T(s)$ is the asset payoff in state s. In consumption-based asset pricing models, the pricing kernel is the marginal rate of substitution of state-dependent consumption at dates t and T, and many papers have estimated the pricing kernel using variations of equation (1) as an identifying condition. It is typically assumed that the underlying price distribution is stationary. An unconditional version of (1) is then used.

For derivative assets with payoff function $g(S_T(s))$, the current derivative price D_t is:

(2) $$D_t = E_t\left[K_{t,T}(s)g(S_T(s))\right]$$

Equation (2) indicates that derivatives prices provide a rich set of additional identifying conditions for the pricing kernel. It is clear that knowledge of the pricing kernel and the objective conditional probability measure is sufficient to price any derivative asset with payoffs dependent only on the underlying asset price at time T.

Equation (2) may be rewritten to emphasize the two objects which are estimated in this paper: the pricing kernel $K_{t,T}(s)$ and the objective conditional probability measure $f_{t,T}(s)$.

(3) $$D_t = \int K_{t,T}(s)g(S_T(s))f_t(S_T(s))dS_T(s)$$

This paper uses the fact that the pricing kernel is interpreted as the state price per unit probability. Thus, the ratio of estimated state prices and state probabilities may be used as an estimate of the pricing kernel. For estimation purposes, it is useful to discretize the state space so that states are defined by discrete return ranges. The state prices are then interpreted as the prices of the supershares of Hakansson (1976) or the delta securities of Breeden and Litzenberger (1978).

In this case, the states used are 100 equiprobable ranges (i=1...100) based on a lognormal diffusion with fixed parameters and fixed time interval. Let s_i be the subset of the real line corresponding to the return range for state i. With the state space divided into 100 return states, non-parametric estimation of the EPK may be accomplished by taking the ratio of the estimated discrete state price for state i, $z_{t,T}(s_i)$, with the estimated discrete probability of state i, $p_{t,T}(s_i)$.

(4) $$K_{t,T}(s_i) = \frac{z_{t,T}(s_i)}{p_{t,T}(s_i)}$$

Equation (4) provides a non-parametric estimate of the EPK in the sense that the functional form of the EPK is unrestricted, although restrictions on state prices and probabilities are reflected in the EPK. Notice that this specification allows the pricing kernel to implicitly depend on many economic variables such as current and lagged aggregate consumption, the level of interest rates, and so forth.

II.b. Estimating state prices

The first step in EPK estimation involves the state price density, $\psi_{t,T}(r_{t,T})$, and the discrete state prices, $z_{t,T}(s_i)$. The continuous state price density is defined over net return states, $r_{t,T}$, for the period from t until T, where the conditioning variables are suppressed. The state price density is estimated using the flexible risk-neutral density function method of Rosenberg (1995).

Consider the following state-price and risk-neutral density formulation of equation (3). Letting the payoff function g(•) depend on the T-t period net return and suppressing dependence on S_t:

(5) $$D_t = \int \psi_{t,T}(r_{t,T}) g(r_{t,T}) dS_T = e^{-r(T-t)} \int g(r_{t,T}) f^*_{t,T}(r_{t,T}) dr_T$$

Equation (5) provides identifying conditions so that the state price density may be estimated from a cross-section of derivatives prices. Normalized state prices, also known as the risk-neutral probabilities defined by the risk-neutral density $f^*_{t,T}(r_{t,T})$, are estimated based on a flexible density function which allows for a variety of tail shapes, but incorporates the lognormal as a special case. The flexible density function used in this paper is a modified version Rosenberg (1995) formulation, so that the null model for the risk-neutral density of returns is the lognormal density implied by a diffusion process over a period of length T-t with drift rate μ and diffusion parameter σ.

The flexible density generalizes the lognormal density by allowing the σ parameter to be a function of the level of the state variable. However, this is not the density implied by a diffusion with drift rate dependent on the level. The intuition behind the flexible density function is that, at a given point, the actual probability may deviate from the lognormal probability with constant variance. The actual probability at the point may be matched by adjusting the variance of the lognormal distribution as needed.

The function that maps the lognormal probability into the actual probability is called the sigma shape polynomial (SSP). The SSP and its graph concisely summarize the deviation of the risk-neutral density from the lognormal. Stochastic volatility effects should result in higher tail probabilities consistent with lognormal distributions with higher values of σ. In this case, the sigma shape polynomial would be an upward parabola. Leverage effects may result in left skewness causing a negatively sloped SSP which puts greater probabilities on large negative outcomes than for a flat SSP.

Equations (6), (7), and (8) specify the flexible density function used in this paper. Equation (6) defines the flexible density function with scaling parameter given by equation (7). Equation (8) defines the sigma shape polynomial, which in this case is log-quadratic in log-gross returns.

(6) $$f^*_{t,T}(r_{t,T};\lambda_t,\mu_t,\theta_t) = \text{k} \exp\left[-.5\left(\frac{\log(r_{t,T}+1)-(\mu_t-.5\sigma(r_{t,T};\theta_t)^2(T-t))}{\sigma(r_{t,T};\theta_t)\sqrt{T-t}}\right)^2\right]$$

(7) $$k = \lambda_t \frac{1}{X_T \sqrt{2\pi\sigma(r_{t,T};\theta_t)(T-t)}}$$

(8) $$\sigma(r_{t,T};\theta_t) = \exp\left[\beta_{1,t} + \beta_{2,t}\log(r_{t,T}+1) + \beta_{3,t}\log(r_{t,T}+1)^2\right]$$

Notice that this particular specification defines a density family with the five parameters λ_t, μ_t, $\beta_{1,t}$, $\beta_{2,t}$, $\beta_{3,t}$. λ_t is a scaling factor necessary to ensure integration to unity, and the density function is defined over a large but bounded support. $\beta_{2,t}$ and $\beta_{3,t}$ are the terms that allow the sigma shape polynomial to deviate from lognormality. The parameters of $f^*_{t,T}(r_{t,T})$ are time-varying reflecting the fact that time-varying objective probabilities generate time variation in state prices not captured by Black-Scholes.

Using the return boundaries for the 100 prespecified states, the continuous state price density is mapped to the discrete state prices using numerical integration.

(9) $$z_{t,T}(s_i) = \int_{r_{t,T}\in s_i} \psi(r_{t,T})dr_{t,T} = e^{-r_f(T-t)} \int_{r_{t,T}\in s_i} f^*(r_{t,T})dr_{t,T} \qquad i=1..100$$

II.c. Estimating objective state probabilities

The second stage in EPK estimation involves the time-varying objective probability density, $f_{t,T}(r_{t,T})$, and the discrete state probabilities $p_{t,T}(s_i)$. One of the fundamental forms of time-variation in the objective density, especially over short time horizons, is stochastic volatility. In this paper, the GARCH components with leverage (GCOMP) model (Engle and Lee, 1993) is used to provide a parsimonious description of the time-varying conditional volatility process. The GARCH components model allows for a long and short run component of volatility that decay at different rates, and an asymmetric effect of news on volatility.

The model is specified as follows:

(10) $$\ln(S_t / S_{t-1}) = \mu + \varepsilon_t,\ \varepsilon_t \sim N(0,\sigma_t^2)$$

(11) $$\sigma_t^2 = q_t^2 + \alpha(\varepsilon_{t-1}^2 - q_{t-1}^2) + \gamma(Max[0,-\varepsilon_{t-1}]^2 - .5q_{t-1}^2) + \beta(\sigma_{t-1}^2 - q_{t-1}^2)$$

(12) $$q_t^2 = \omega + \rho q_{t-1}^2 + \phi(\varepsilon_{t-1}^2 - \sigma_{t-1}^2)$$

While it is clear that the constant risk premium formulation as in equation (10) is not compatible with time-varying state prices, the effect over the short time horizon considered in this paper is not likely to be important, and equation (10) may be viewed as an approximation. Even if this particular volatility model is misspecified, GARCH models often exhibit useful approximating characteristics for other volatility processes. In addition, to reduce the potential effects of model misspecification, the empirical standardized residuals are saved from the estimation procedure and used when the state probabilities are estimated. The GCOMP model is also re-estimated each day using current data, which will capture potential non-stationarity in the parameters.

The conditional discrete state probabilities on date t are estimated by simulating N=100,000 returns based on an estimated GCOMP model and calculating the observed probability of a return within each state's return range.

(13) $$p_{t,T}(s_i) = \frac{1}{N}\sum_{j=1}^{N} I_{r_{t,T,j} \in s_i} \quad I(\bullet) \text{ is an indicator function} \qquad i = 1..100$$

II.d. Estimating EPK hedge ratios

This paper introduces an option hedging technique that utilizes the empirical pricing kernel. This hedging technique offers two primary advantages over existing methods. First, option hedging using an EPK incorporates current information about investor preferences and forecast probabilities. This should improve hedge performance in markets characterized by time-variation in either of these factors. Second, option hedging using an EPK directly solves the problem of minimum variance hedging rather than relying on a continuous approximation. The EPK hedging methodology could also be applied to hedging downside risk or value-at-risk.

Suppose the objective of the hedger is to minimize hedge portfolio variance for a one-day hedge. In particular, let D_1 be the price of the instrument being hedged, let D_2 be the price of another derivative asset, and let S be the price of the underlying asset. The hedger chooses the optimal holdings of the two additional assets by solving the program:

(14) $$\underset{a_{1,t},\, a_{2,t}}{Min}\ Var_t\left[(D_{1,t+1} - D_{1,t}) + a_{1,t}(D_{2,t+1} - D_{2,t}) + a_{2,t}(S_{t+1} - S_t)\right]$$

In general, the hedge portfolio composition will change through time reflecting changes in conditional probabilities, the EPK, and the moneyness of the options. Equation (14) can be applied to the single asset hedge by constraining $a_{1,t}$ or $a_{2,t}$ to be zero.

The difficulty in solving equation (14) is that, as of date t, the derivative prices and underlying price at t+1 are distributions which may not have closed form representations. To surmount this problem, a simulation technique is used. By holding the pricing kernel at date t constant until t+1, the joint distribution of prices at t+1 may be obtained.

The use of the date t pricing kernel for pricing at date t+1 might be viewed as a form of martingale assumption. Notice that risk due to the stochastic behavior of the pricing kernel over one day is not considered in this formulation, but it will be addressed in a companion paper.

Initially, 100 simulated underlying prices at t+1 ($S_{t+1,j}$; j=1...100) are generated by sampling independent draws with replacement of innovations from an empirical distribution and using the one day ahead GCOMP conditional variance forecast, $\sigma^2_{t+1|t}$. The empirical distribution is defined as the distribution of standardized GCOMP residuals from a model estimated over a fixed time period.

$$(15) \quad \tilde{S}_{t+1,j} = S_t \exp^{\mu_t + \sqrt{\sigma^2_{t+1|t}} * e_j} \qquad e_j \sim D_{empirical} \quad j=1..100$$

The fundamental pricing equation (2) may be then used to approximately price derivatives at date t+1 by simulating underlying price returns from t+1 until T and substituting the date t pricing kernel for the date t+1 pricing kernel. Each simulated return is mapped to one of the one-hundred possible discrete state returns based on the state return range.

Define the discrete state return for state i, $r_{i,t+1,T,j}$, as the midpoint of each state return range for states 2 through 99 and as the lower and upper return range bound for states 1 and 100 respectively. Let the payoff function for the derivative with price D_1 be $g_1(\bullet)$ with dependence on $S_{t+1,j}$ suppressed. Then, the j[th] simulated derivative price t+1 is:

$$(16) \quad D_{1,t+1,j} = \frac{1}{1000} \sum_{k=1}^{5000} \sum_{i=1}^{100} K_{t,T}(s_i) I_{r_{t+1,T,j} \in s_i} g_1(r_{i,t+1,T,j})$$

Equation (16) may be used to price other derivative assets by substituting the appropriate payoff function, e.g. $g_2(\bullet)$. The assumption of a constant pricing kernel over one day is evident in this equation.

This procedure generates 100 realizations of underlying and derivative prices at t+1. The minimum variance portfolio weights are obtained by regression of the price changes of the instrument to be hedged on price changes of the other two instruments. This appears similar to minimum variance hedging of futures, but in this case, the regression is not over time, but over state realizations at t+1.

III.a. Estimation of the S&P500 state prices and probabilities

S&P500 state prices and probabilities are estimated on 839 of 1243 days in the sample period based on data availability for the hedging tests. The criteria for inclusion in the hedging tests is described in section IV.

Table 1 summarizes the characteristics of the S&P500 index option data provided by the CBOE. State price estimation utilizes a cross-section of all non-in-the-money call and put option prices on each hedging date with the same maturity as the selected ATM and OTM options. The second panel of Table 1 shows that on average, there are 16 non-ITM options available per day with a moneyness range of 11%. Moneyness is defined as S/K-1 for calls and K/S-1 for puts, where S is the current underlying price and K is the option exercise price.

Figures 1 and 2 summarize the state price estimation results. Figure 1 plots the average sigma shape polynomial over the sample period. The downward sloping curve reflects the volatility skew often found in post-crash equity option data and suggested by the ATM and OTM implied volatilities. The skew implies that state prices for large negative return states are relatively high compared to those of the Black-Scholes model. At this stage, it cannot be determined whether this pattern is due to higher probabilities associated with these states or increased investor preferences for positive payoffs in these states. Figure 2 plots the daily sigma shape polynomials estimated over the period May through June 1996. It is clear that the state prices vary from day to day, although the general shape of the SSP is consistent.

State probability estimation is based on historical S&P500 index return data. The sample properties of the S&P500 index return over the hedging period are described in the first panel of Table 2. S&P500 return volatility was fairly low with a return standard deviation of .57% per day. This is substantially below the average ATM or OTM implied standard deviation. S&P500 returns exhibited negative skewness and positive kurtosis, which is consistent with a stochastic volatility model with asymmetric effects. The average daily S&P500 index return was .06%.

Table 2 also describes a GCOMP model estimated over the period 1982-1996, a period substantially longer than the hedging period. The standardized residuals from this model define the empirical distribution of the innovations in the state probability and EPK hedge ratio estimation. The leverage term, which indicates that large negative returns disproportionately increase volatility, is significant using a standard t-statistic but only marginally significant using a robust t-statistic. The long run and short run volatility components are both significant, with a highly persistent long run component indicated by a value of ρ close to 1. The standardized residuals deviate significantly from normality as illustrated by their negative skewness and positive excess kurtosis.

The last panel of Table 2 reports the iteratively estimated out-of-sample GCOMP models used in the daily estimation of the state probabilities. The initial estimation is performed using 1000 lagged returns, and subsequent estimations use an expanding window of returns. The average estimates are somewhat different than the estimates in the long sample estimation, suggesting possible non-stationarity in the parameters over the sample period. The average iterated leverage estimate is smaller than the full sample leverage estimate, and average iterated long run volatility persistence is lower than the full sample estimate.

Figure 3 plots the GARCH daily average conditional volatility for a T-t period horizon, the daily implied volatilities for a T-t period horizon, and 20-day rolling standard deviations over the sample period. The GCOMP estimates and ATM implied volatility estimates track most closely, although the ATM implieds are more volatile than the GCOMP forecasts. The OTM implieds are consistently the highest volatilities, and the 20 day standard deviations are lower than the other volatilities over most of the period.

III.b. Estimation of the S&P500 empirical pricing kernel

Using the methodology described in section II and the results from section III.a, the S&P500 empirical pricing kernel is estimated on a daily basis over the period 1992 through 1996. Initially, it is useful derive the benchmark pricing kernel implied by the Black-Scholes model for comparison to the EPK.

The pricing kernel implied by the Black-Scholes model is simply the ratio of the BS state prices as derived by Breeden and Litzenberger (1978) and the state probabilities implied by a diffusion with parameters set equal to the sample values over a fixed period. In this paper, the coefficient of relative risk aversion is defined as the opposite of the exponent in the power pricing kernel function. The BS coefficient of relative risk aversion may be derived by taking the ratio of the state price density and the objective probabilities and writing the solution as the product of a constant and a power function.

(17) $$\gamma_{BS} = \frac{\mu - r}{\sigma^2}$$

Using the average one-month Treasury bill rate as the riskless rate and sample moments of daily returns for the periods 1992-1996 and 1962-1996, the implied BS risk aversion coefficients are 14.7 and 5.7, respectively. These levels of risk aversion are substantially higher than obtained from GMM estimation using aggregate consumption data. For example, Hansen and Singleton (1982, 1983) find risk aversion coefficients less than 2.

The implied BS risk aversion coefficients illustrate the equity premium puzzle of Mehra and Prescott (1985). There is a divergence between the empirical risk-premium which reflects investor preferences and the implications of many consumption-based representative agent utility models with plausible levels of risk aversion. Cochrane and Hansen (1992) utilize the Hansen and Jagannathan (1991) framework to examine the compatibility of several utility specifications with historical return moments. They find that risk aversion coefficients of at least 40 are required for CRRA utility, and coefficients of at least 7.5 are required when habit persistence is incorporated. Thus, the high BS risk aversion coefficients are consistent with previous results.

At this point, it is useful to find a single summary measure of EPK risk aversion that may be directly compared with the BS risk aversion coefficient. One way to summarize the EPK is via a power function with a single shape parameter γ_t. This parameter summarizes investor preferences on date t for payoffs over the states on date T and may be interpreted as a local coefficient of relative risk aversion. The following optimization program, which minimizes the squared distance between the realized and fitted EPK, is used for estimation of γ_t.

(18) $$\underset{\gamma_t}{Min} \sum_{i=2}^{99} \left[K_{t,T}(s_i) - (1 + r_i)^{-\gamma_t} \right]^2$$

In this case, r_i is discrete state return which is defined as the midpoint of the state return range for states 2 through 99. States 1 and 100 are defined over return ranges substantially larger than the other states, so they are excluded from this estimation. State 1, which might be viewed as the crash state, is of direct interest, and its time series behavior is analyzed separately.

The average EPK fitted γ is 14.8, and the range is from 1.4 to 30.2. The average EPK γ is quite close to the BS γ for the sample period of 1992-1996, but substantially higher than the BS γ for the longer time period 1962-1996. In addition, there is a large amount of day-to-day variation in preferences as measured by the EPK. Strong persistence in the daily EPK γ is illustrated by a first order autocorrelation of .72.

The closeness of the estimated BS and EPK γ obscures important differences in the two functions. Figure 4 plots the average empirical pricing kernel over the sample period along with the two pricing kernels implied by the Black-Scholes model using the two sets of sample moments, a time horizon of 20 days, and the average one-month Treasury bill rate from 1992-1996. EPK preferences indicate that payoffs in the largest negative return state are more desirable relative to implied BS preferences.The average probability standardized value to an investor of a $1 payoff in a -9% S&P500 return state is about $5.90 using the EPK and about $5.30 or $1.90 using the BS pricing kernels. EPK preferences also suggest that payoffs in large positive return states are less desirable relative to implied BS preferences.

Figure 5 illustrates the day-to-day variability in the EPK over the period May through June 1996. The daily EPK shapes are similar, but the level of risk aversion varies substantially. Figure 6 plots time-variation in the average annual EPK, while figure 7 plots the average EPK and the fitted average EPK using a power function. By fitting a power function to the average EPK, a γ of 15.6 is obtained, which is slightly higher than the average daily EPK γ of 14.8.

Figure 8 illustrates the difference between conditional (EPK) and unconditional (BS) preference models. In this figure, the time-series of EPK γ's is plotted along with the BS γ for the sample period. The BS γ will only change when either the drift, volatility, or riskless rate changes. Of course, a maintained assumption in the Black-Scholes setting is that all of these parameters are constant over time. In contrast, the EPK γ exhibits a substantial amount of variability due to changes in state prices and probabilities which suggests that investor preferences are stochastic.

Figure 9 plots the empirical and implied BS preferences over the crash state, i.e. a T-t day return less than -9%. The BS implied preferences are derived using S&P return moments from 1992 through 1996, the average riskless rate over this period, and a time-horizon of 20 days. While the EPK suggests greater average preference for payoffs in the crash state than BS, the more dramatic difference is the variation in EPK crash preferences compared to BS. The crash state price per unit probability ranges from $.02 to $38.30 for the EPK.

The first panel of Table 3 contains the results of a regression of the daily EPK γ on its first lag and several additional variables. The purpose of this analysis is two-fold. First, if the daily EPK γ is considered to be an adequate representation of investor preferences, then the sources of time variation in preferences may be identified using projection of γ onto explanatory variables. Conversely, if the relationship between γ and other explanatory variables is plausible, then this suggests that γ is a useful measure of investor preferences.

About 63% of the variation in γ is explained in the regression. Most significantly, the γ regression results indicate that the level of risk aversion is negatively correlated with changes in the S&P500 index. In other words, a drop in the S&P500 index level is related to an increase in risk aversion. In particular, a one-day 5% drop in the S&P500 index is associated with an increase of 4.1 in the level of risk aversion. This may also be interpreted as evidence that relative risk aversion is declining in the level of wealth.

The regression results also show that the volatility spread (ATM-GCOMP) is positively correlated with the level of risk aversion. This result is intuitive and expected, since this spread can be considered as a proxy for the risk-premium based on how much the market is pricing options above the objective volatility. Even in the presence of other variables, γ is strongly autocorrelated with a coefficient on it's first lag of .65.

The change in time to expiration variable indicates that as the time to expiration is declining, γ declines. The lagged EPK for state 1 and the change in interest rates are not significant in the regression. Table 3 provides similar results for the time-series behavior of preferences over the crash state, $K_{t,T}(s_1)$, which is abbreviated as k1(t).

IV. Hedging tests

The EPK hedging methodology is applied to the problem of hedging out-of-the-money (OTM) S&P500 index put options using at-the-money (ATM) put options and the index with the objective of minimizing the variance of one-day hedge portfolio price changes. This particular application is informative since writing OTM S&P500 options has been a historically profitable strategy, but methods of effectively hedging this position are not well known.

For the EPK hedge, S&P500 state prices and state probabilities are estimated on a daily basis over the period 1992 through 1996. As described in section III, state prices are derived from S&P500 index option premia and the flexible risk-neutral density function method (Rosenberg, 1995), while probabilities are obtained using simulation of an estimated GARCH components with leverage model

(Engle and Lee, 1993) for S&P500 index returns. The EPK for S&P500 states is the ratio of state prices and the state probabilities.

The EPK hedge portfolio weights are estimated daily according to the methodology in section II.d. Three separate regressions are run to estimate the minimum variance EPK hedge portfolios. These are: a regression of the simulated one-day OTM price change on the simulated underlying price change, a regression of the simulated OTM price change on the simulated ATM price change, and a regression of the simulated OTM price change on both price changes.

In addition to the EPK hedge method, there are four alternatives analyzed for hedging using only the underlying, only the ATM option, and both assets. The first hedging method used is Black-Scholes delta and delta-gamma hedging. In this case, a 20-day historical standard deviation is used to calculate the pricing sensitivities.

An alternative to Black-Scholes delta-gamma hedging is considered to ensure that the Black-Scholes model is not rejected because the hedging occurs in discrete-time. This method is the Black-Scholes minimum variance hedge. In this case, a simulation is used to generate 5000 realizations of one day ahead underlying and ATM option prices. The underlying price process is specified using the fixed historical S&P500 moments over sample period. The minimum variance portfolio weights are estimated by a regression of OTM price changes on the price changes of the hedging instruments.

A historical regression hedge is also used and is constructed by a rolling 40 day regression of OTM price changes on price changes of the hedging instrument. The historical hedge portfolio weights are updated daily.

The hedging method which is intended to represent local volatility function (LVF) models is the LVF hedge. The LVF technique used in this paper is somewhat different than deterministic volatility function (DVF) methods in the existing literature. Usually in DVF models, the sigma function (or the implied risk-neutral probability density) is estimated as a deterministic function of the level of the underlying price and time using the current cross-section of option prices. The theoretical hedge consistent with the DVF model should be based on the dynamics implied by a fixed DVF. However, in practice, the instability of the DVF makes it is clear that the effects of changes in the DVF must be hedged as well as changes due to the underlying price and time.

The primary feature of the DVF techniques is their ability to match a current cross-section of option prices. In this paper, this feature is modeled in a LVF by allowing each option to be priced using the Black-Scholes formula with an option specific implied volatility. As an approximation to the option price dynamics, it is assumed that option sensitivities to underlying price changes are given by the Black-Scholes delta and gamma evaluated at the option-specific implied volatility. One justification for this method is that local price movements might be approximated by a local version of the Black-Scholes pricing formula. This technique results in substantially different hedge portfolio weights than Black-Scholes hedging.

The LVF method, unlike the EPK method, reveals nothing about the relative effects of time-varying preferences versus time-varying probabilities or the effects of non-diffusion price dynamics on hedge ratios. However, the LVF method does capture the joint departure of probabilities and preferences from Black-Scholes at an instant in time.

The first panel of Table 1 reports characteristics of the ATM and OTM put options used in the hedging tests. On each day, the nearest maturity put options with at least 10 but no more than 60 trading days until expiration are chosen. Then, the ATM option is selected as the nearest to the money put option, which must be within 1% of the money. The OTM option is selected as the option which is closest to, but at least, 3% out-of-the-money. The average ATM moneyness is near 0%, the average OTM moneyness is -3.5%, and the average time to expiration is approximately 20 trading days.

The average ATM and OTM implied volatilities of .74% and .94% suggest that the period is characterized by an implied volatility skew. That is, OTM implied volatilities are higher than ATM implied volatilities. The standard deviations of one-day ATM and OTM implied volatility changes illustrate that implied volatilities are not constant over the sample period.

Table 4 reports the hedging test results, including both the standard deviation and interquartile range of hedge portfolio price changes. While the standard deviation of price changes was chosen as the criterion to minimize, several outliers in the raw option price change data suggest that a more robust variability measure should be reported as well.

The results of using the five methods to hedge a written \$100 OTM put position using only the underlying asset are quite interesting. The unhedged OTM put position has a standard deviation of price changes equal to \$24.60 per day. The most effective hedge in this case is the LVF delta hedge which reduces the portfolio standard deviation by about 40% to \$14.90 per day. The least effective hedge is the BS delta hedge which reduces the standard deviation by about 18% to \$20.16 per day. The EPK hedge is the third most effective hedge with a standard deviation of \$15.10, but is the most effective hedge based on IQR. All of the hedge methods require a fractional short position in the underlying asset.

The next hedge considered is a hedge of the OTM option using only the ATM option, which is a form of a ratio spread. Since the ATM option is a more similar asset to the OTM option than the underlying, it is not surprising that all methods offer some improvement over the hedging using the underlying. Again, the most effective hedge is the LVF hedge which reduces the portfolio standard deviation by about 50% to $12.58. The EPK method comes close to this performance with a hedge portfolio standard deviation of $12.80. Again, the EPK method is most effective hedge based on IQR. All of the hedge methods require a fractional long position in the ATM option.

The final hedge considered uses both the ATM option and the underlying as hedging instruments. Since their price changes are not perfectly correlated, one would expect that a two asset hedge would improve on a one asset hedge. In fact, only two methods show a slight improvement, and they both are inferior to the LVF and EPK methods. In this case, the EPK method is most effective based on both the standard deviation and IQR criteria.

Several other results are as follows. The BS minimum variance hedge is consistently more effective than the standard BS hedge. However, the BS minimum variance hedge has the advantage of using the realized S&P moments over the sample period. In addition, the historical regression hedge is a significant improvement over the BS hedges.

Overall, the hedge results indicate that the EPK method is a superior hedge technique based on the standard deviation criterion relative to the alternatives except the LVF method. When the more robust IQR criterion is used, the EPK method is superior to the LVF method as well.

V. Conclusions

This paper develops an option hedging methodology based on daily estimation of the empirical pricing kernel (EPK) or state price per unit probability. This measure of investor preferences over states is linked to a dynamic model of the underlying price to obtain a one-day ahead forecast of derivative price distributions and minimum variance hedge ratios.

The EPK is estimated on a daily basis over the period 1992 through 1996 for states defined by S&P500 returns. The non-parametric EPK estimation technique used does not restrict the functional form of the EPK. State prices are estimated using S&P500 index option premia and the flexible risk-neutral density function method (Rosenberg, 1995), while probabilities are obtained using simulation of an estimated GARCH components with leverage model (Engle and Lee, 1993).

Empirical results indicate that average EPK risk aversion is higher than that implied by the Black-Scholes model using long-run estimates of the mean and variance of S&P500 returns, but is close to implied BS risk aversion using the mean and variance of S&P500 returns over the sample period. In addition, evidence from the estimated daily EPK suggests that investor preferences vary significantly over time. It is found that the daily level of risk aversion is strongly positively autocorrelated, negatively correlated with S&P500 price changes, and positively correlated with the spread between implied and objective volatilities.

Deviations from the Black-Scholes preference and probability assumptions suggest that hedging performance may be improved with a model that better reflects empirical data including the time-varying nature of both preferences and probabilities. Hedging results reveal that typical hedging techniques for out-of-the-money S&P500 index options, such as Black-Scholes or historical minimum-variance hedging, are inferior to the EPK hedging method. The similarity of the LVF hedge performance to EPK performance suggests that in some cases the LVF technique may offer a reasonable approximation of the EPK hedge ratios.

Bibliography

Amin, K. I. and V. K. Ng, 1993, "Option Valuation with Systematic Stochastic Volatility," *Journal of Finance*, 48, 881-910.

Black, F. and M. Scholes, 1973, "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81, 637-654.

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

Brennan, M., 1979, "The Pricing of Contingent Claims in Discrete-time Models," *Journal of Finance*, 34, 53-68.

Cochrane, J. H. and L. P. Hansen, 1992, "Asset Pricing Explorations for Macroeconomics," *NBER Macroeconomics Bulletin*, 115-165.

Derman, E. and I. Kani, 1994, "Riding on a Smile," *Risk*, 7, 32-38.

Duan, J. C., 1995, "The GARCH Option Pricing Model," *Mathematical Finance*, 5, 13-32.

Dumas, B., J. Fleming and R. E. Whaley, 1996, "Implied Volatility Functions: Empirical Tests," NBER Working Paper No. 5500.

Engle, R. F. and G. Lee, 1993, "A Permanent and Transitory Component Model of Stock Return Volatility," University of California - San Diego, UCSD discussion paper, 92-44R.

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.

Hakansson, N. H., 1976, "The Purchasing Power Fund: A New Kind of Financial Intermediary," *Financial Analysts Journal*, 49-59.

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 K. J. Singleton, 1982, "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models," *Econometrica*, 50, 1269-1286.

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

Jackwerth, J. C., 1996, "Recovering Risk Aversion from Option Prices and Realized Returns."

Mehra, R. and E. C. Prescott, 1985, "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15, 145-161.

Robins, R. P. and B. Schachter, 1994, "An Analysis of the Risk in Discretely Rebalanced Option Hedges and Delta-based Techniques," *Management Science*, 40, 798-808.

Rosenberg, J., 1995, "Pricing Multivariate Contingent Claims Using Estimated Risk-neutral Densities," Stern School of Business, FIN-96-36.

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

Rubinstein, M., 1994, "Implied Binomial Trees," *Journal of Finance*, 49, 771-818.

Stapleton, R. C. and M. G. Subrahmanyam, 1984, "The Valuation of Options when Asset Returns are Generated by a Binomial Process," *Journal of Finance*, 39, 1525-1539.

Table 1 - Summary of option data

Data gathered by the Chicago Board Options Exchange, end of day price quotes, Jan. 1992 - Jun. 1996.
Options for which current and next day's price are available are used in the study.
The out-of-the-money (OTM) options are at least 3 percent OTM.
Prices are in points ($100 units). Moneyness is S/K -1 for calls, K/S - 1 for puts.

	ATM put option	OTM put option
Num. obs.	839	839
Average price	6.09	2.08
Std. price	1.95	1.19
Average price change	-0.22	-0.10
Std. dev. price change	1.25	0.51
Average time to matur.	20.02	20.02
Std. time to matur.	6.60	6.60
Average moneyness	0.000	-0.035
Std. moneyness	0.003	0.004
Average implied std.	0.0074	0.0094
Std. dev. chg. istd.	0.0013	0.0013

	Average	Std.dev.
Moneyness range per day for state price estimation	0.11	0.04
Number of options per day for state price estimation	16.12	2.81

Table 2 - Models of S&P500 daily returns

S&P500 daily total returns over all days in the hedging test period
Jan. 1992 - Dec. 1996

Number obs	Mean	Std. Dev.	Skewness	Excess kurtosis
1265	0.0006	0.0057	-0.27	2.47

GARCH components with leverage model for S&P500 index returns
In-sample estimation
Maximum likelihood estimation with normal as the underlying density.
Week of Oct. 1987 crash down-weighted.
3528 observations, log daily total returns for S&P500 index from Jan. 1982 - Dec.1996.

	Coefficient	Std Error	t-stat	Robust Std Err	Robust t-stat	Ljung-Box(15)
μ	0.0008	0.0001	9.29	0.0001	6.86	5.97
ω	8.00E-09	9.00E-09	0.92	4.00E-08	0.20	
α	0.0420	0.0054	7.7475	0.0159	2.64	
β	0.9181	0.0045	204.94	0.0222	41.39	
γ	0.0424	0.0069	6.13	0.0307	1.38	
ϕ	0.0020	0.0015	1.38	0.0092	0.22	
ρ	0.9990	0.0004	2677.13	0.0011	893.01	

Properties of standardized residuals: et/sqrt(ht)

Number obs	Mean	Std. Dev.	Skewness	Excess kurtosis
3528	0.02	1.05	-0.81	8.55

Out-of-sample estimation
Iterative estimation of GCOMP model with expanding window
Initial estimation is with 1000 observations
839 estimated models on hedging dates from Jan. 1992 -June 1996

	Mean estimate	St.dev. estimates
μ	0.0007	0.00005
ω	5.13E-06	1.86E-06
α	0.0160	0.0063
β	0.9463	0.0929
γ	0.0087	0.0232
ϕ	0.0228	0.0026
ρ	0.9118	0.0300

Table 3 - Analysis of the EPK

Daily EPK summarized by a power function with coefficient gamma
1992-1996, 839 observations on hedging dates

Regression of gamma(t) on explanatory variables

Intercept	**4.53**	0.0001
gamma(t-1)	**0.65**	0.0001
change in time to exp.	**-0.15**	0.0001
change in int. rate	-13151.00	0.5776
S&P500 return (t)	**-81.90**	0.0001
Spread (ATM implied std. - GCOMP std.)	**1072.79**	0.0001
k1(t-1)	-0.01	0.7648
Adjusted R2	0.6318	
Durbin Watson	2.048	

Regression of k1(t) on explanatory variables

Intercept	**0.64**	0.0229
gamma(t-1)	**0.51**	0.0001
change in time to exp.	**-0.08**	0.0001
change in int. rate	-28867.00	0.1275
S&P500 return (t)	**-72.28**	0.0001
Spread (ATM implied std. - GCOMP std.)	**830.23**	0.0001
gamma(t-1)	**0.11**	0.0001
Adjusted R2	0.6291	
Durbin Watson	2.172	

Table 4 - Hedging test results

Options on S&P500 index, Jan. 1992 - Jun. 1996 (839 observations with all data).
Hedging a $100 OTM written put position using other instruments.

Portfolios	Standard deviation of daily price changes	Interquartile range of daily price changes	Average number of units of underlying per written OTM put	Average number of units of ATM option per written OTM put
No hedge:				
100$ OTM written put position	24.60	28.33	0.00	0.00
Hedge using underlying:				
Black-Scholes delta hedge	20.16	21.08	-0.06	0.00
Black-Scholes minimum variance hedge	19.03	20.78	-0.06	0.00
Historical (40 day) regression hedge	14.92	17.51	-0.16	0.00
LVF delta hedge	14.90	17.71	-0.18	0.00
EPK hedge (underlying only)	15.10	16.73	-0.12	0.00
Hedge using ATM put:				
Black-Scholes delta hedge	19.60	20.89	0.00	0.14
Black-Scholes minimum variance hedge	17.86	20.03	0.00	0.15
Historical (40 day) regression hedge	14.25	15.87	0.00	0.33
LVF delta hedge	12.58	13.67	0.00	0.37
EPK hedge (ATM only)	12.80	12.82	0.00	0.30
Hedge using underlying and ATM put:				
Black-Scholes delta-gamma hedge	19.80	21.81	0.07	0.28
Black-Scholes minimum variance hedge	17.82	20.34	0.07	0.30
Historical (40 day) regression hedge	13.99	15.86	-0.06	0.23
LVF delta-gamma hedge	13.90	14.85	0.07	0.51
EPK hedge (underlying and ATM)	13.61	14.39	0.07	0.43

Figure 1
Average sigma shape polynomial
S&P500 index options, 1992-1996

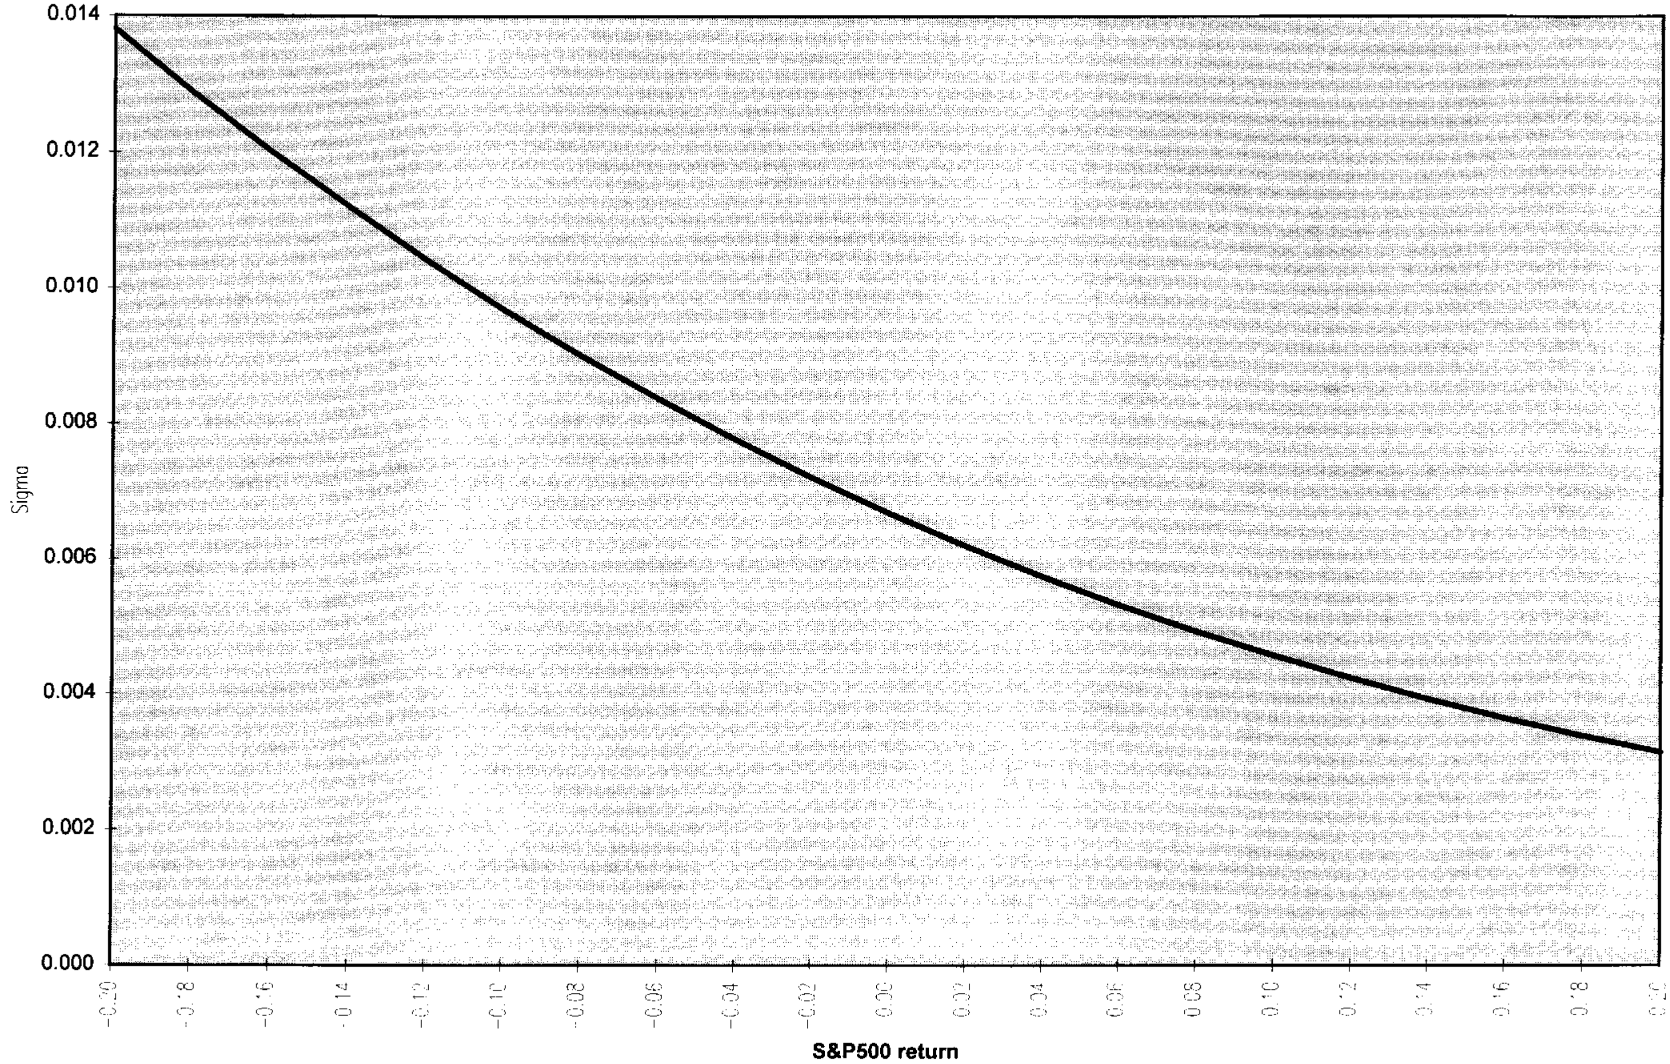

Figure 2
Daily sigma shape polynomials
S&P500 index options, May-June 1996

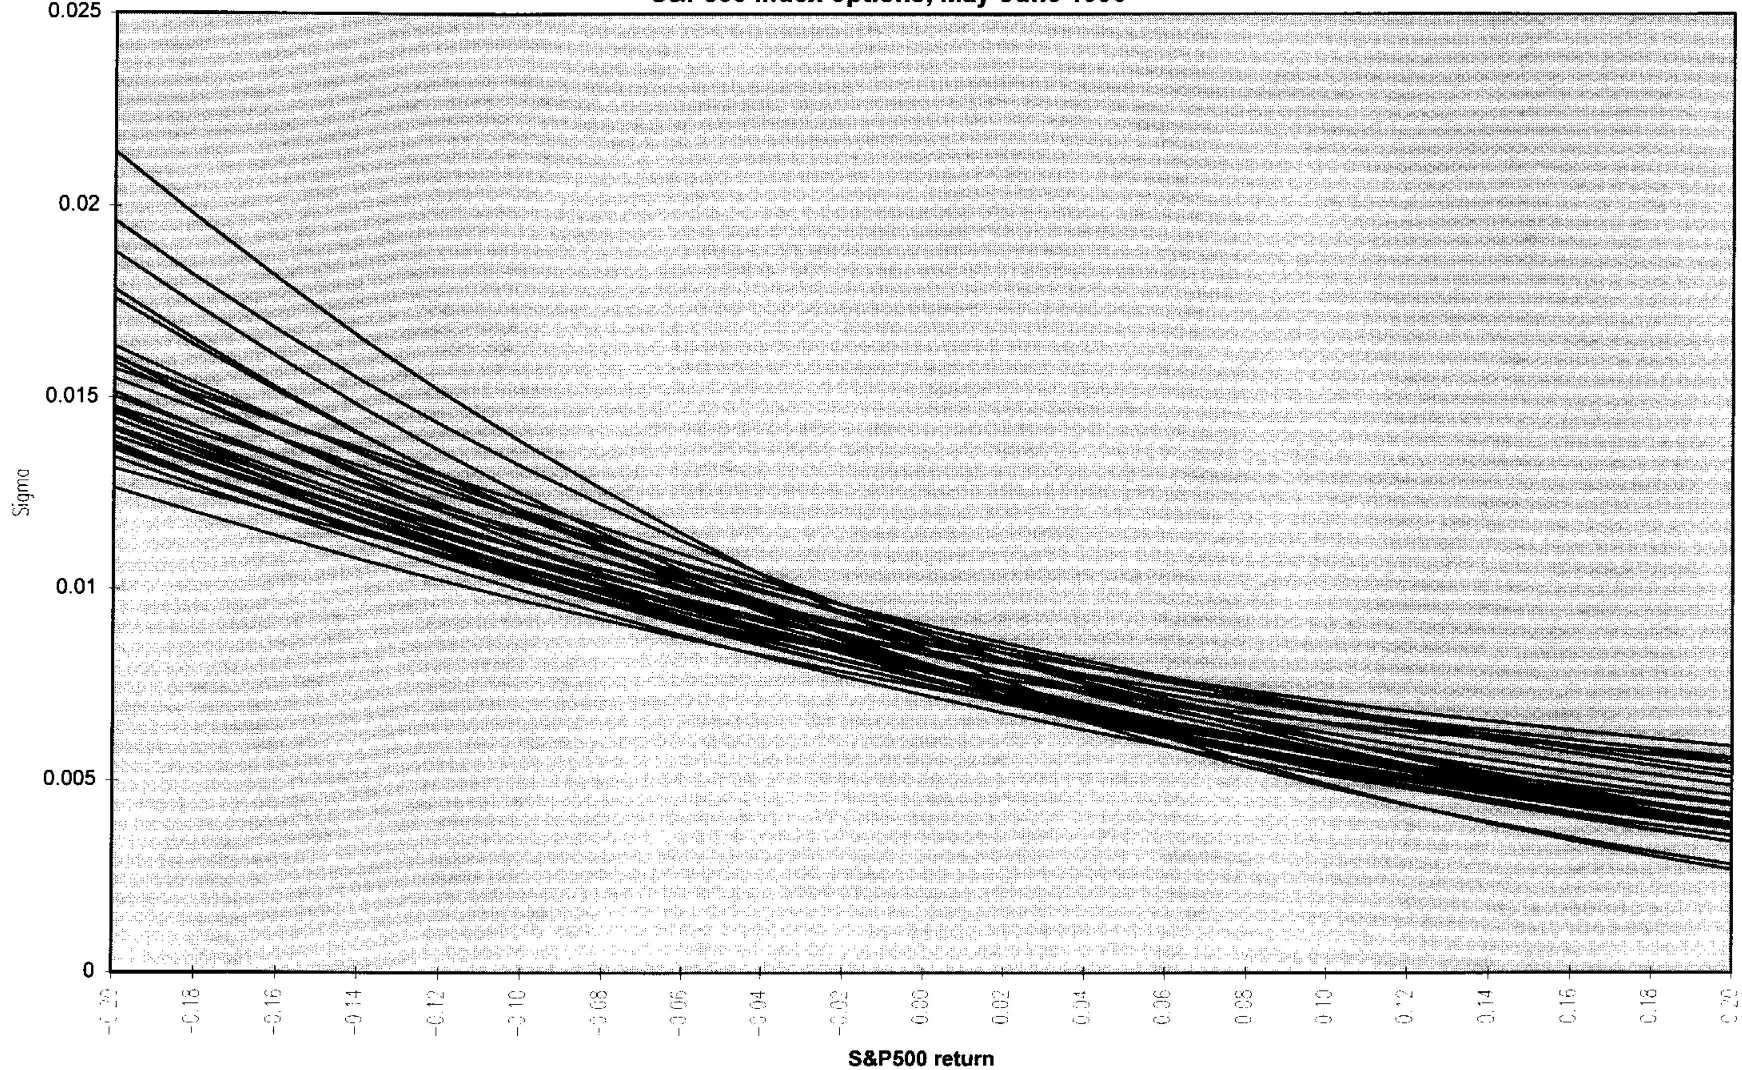

Figure 3
Comparison of volatility estimates
Time horizon (T-t)

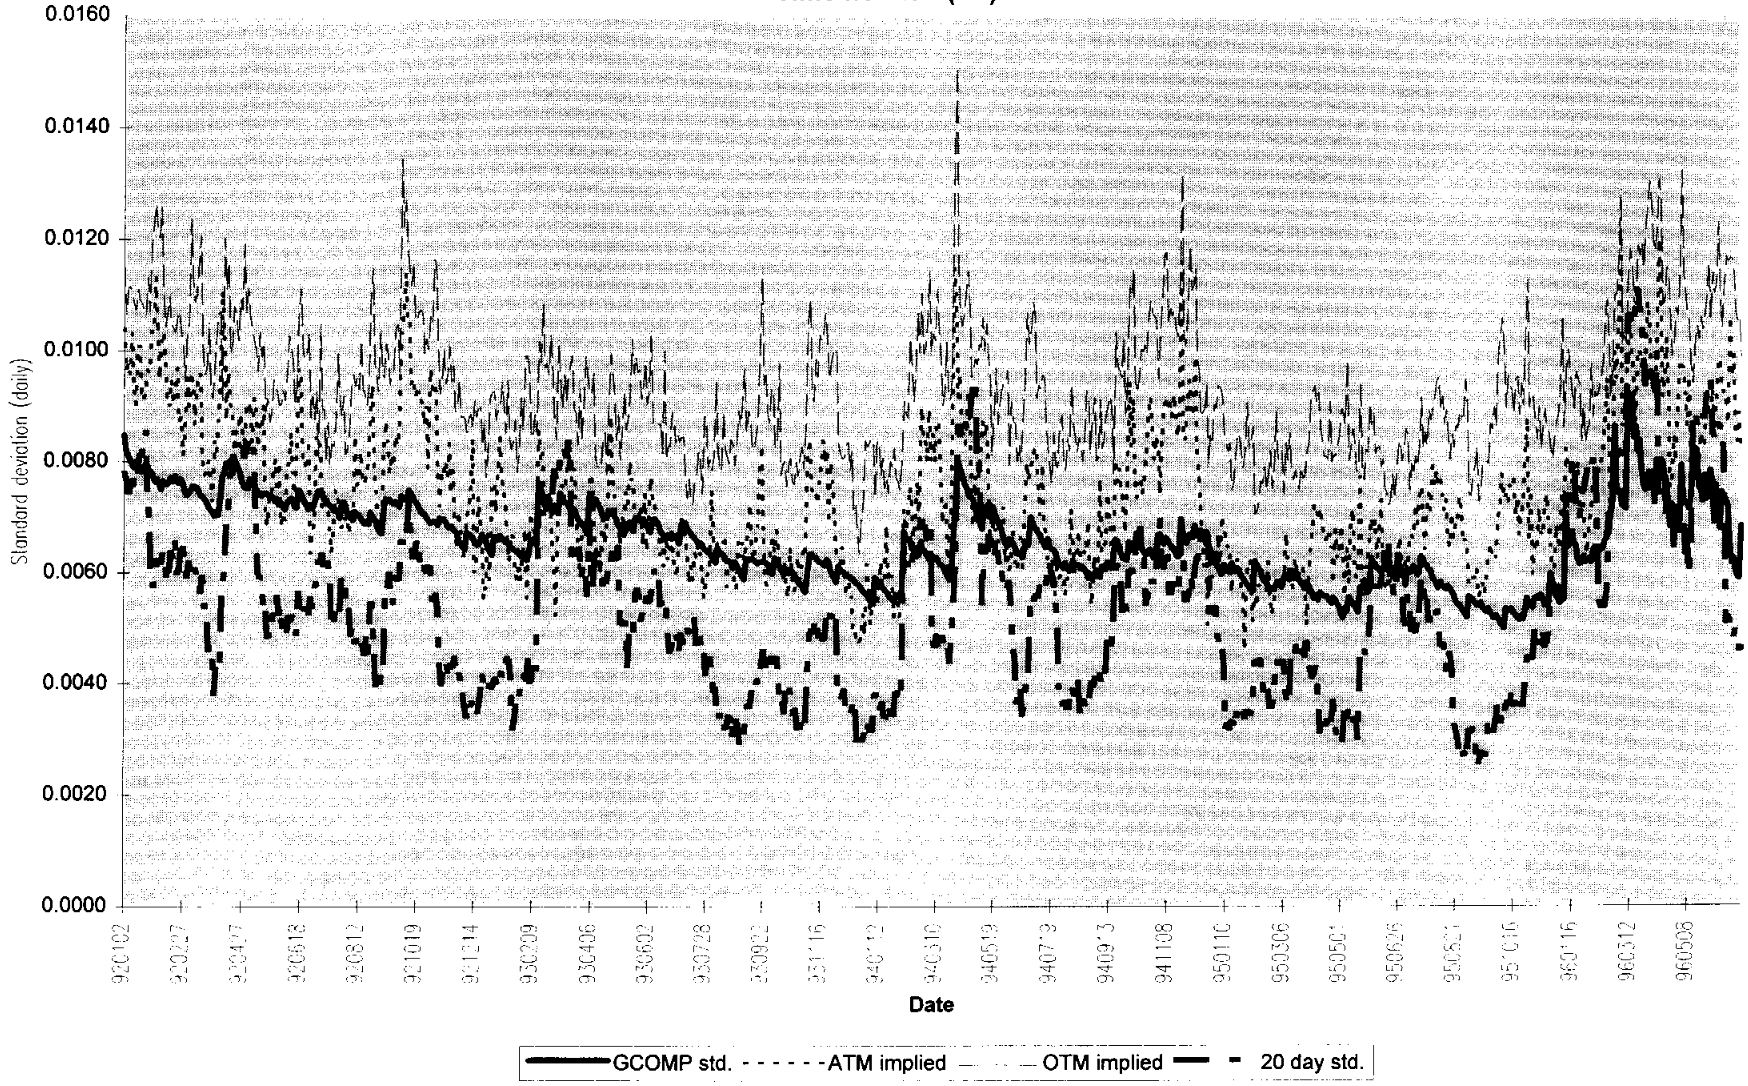

Figure 4
Average pricing kernels for S&P500 return states
1992-1996

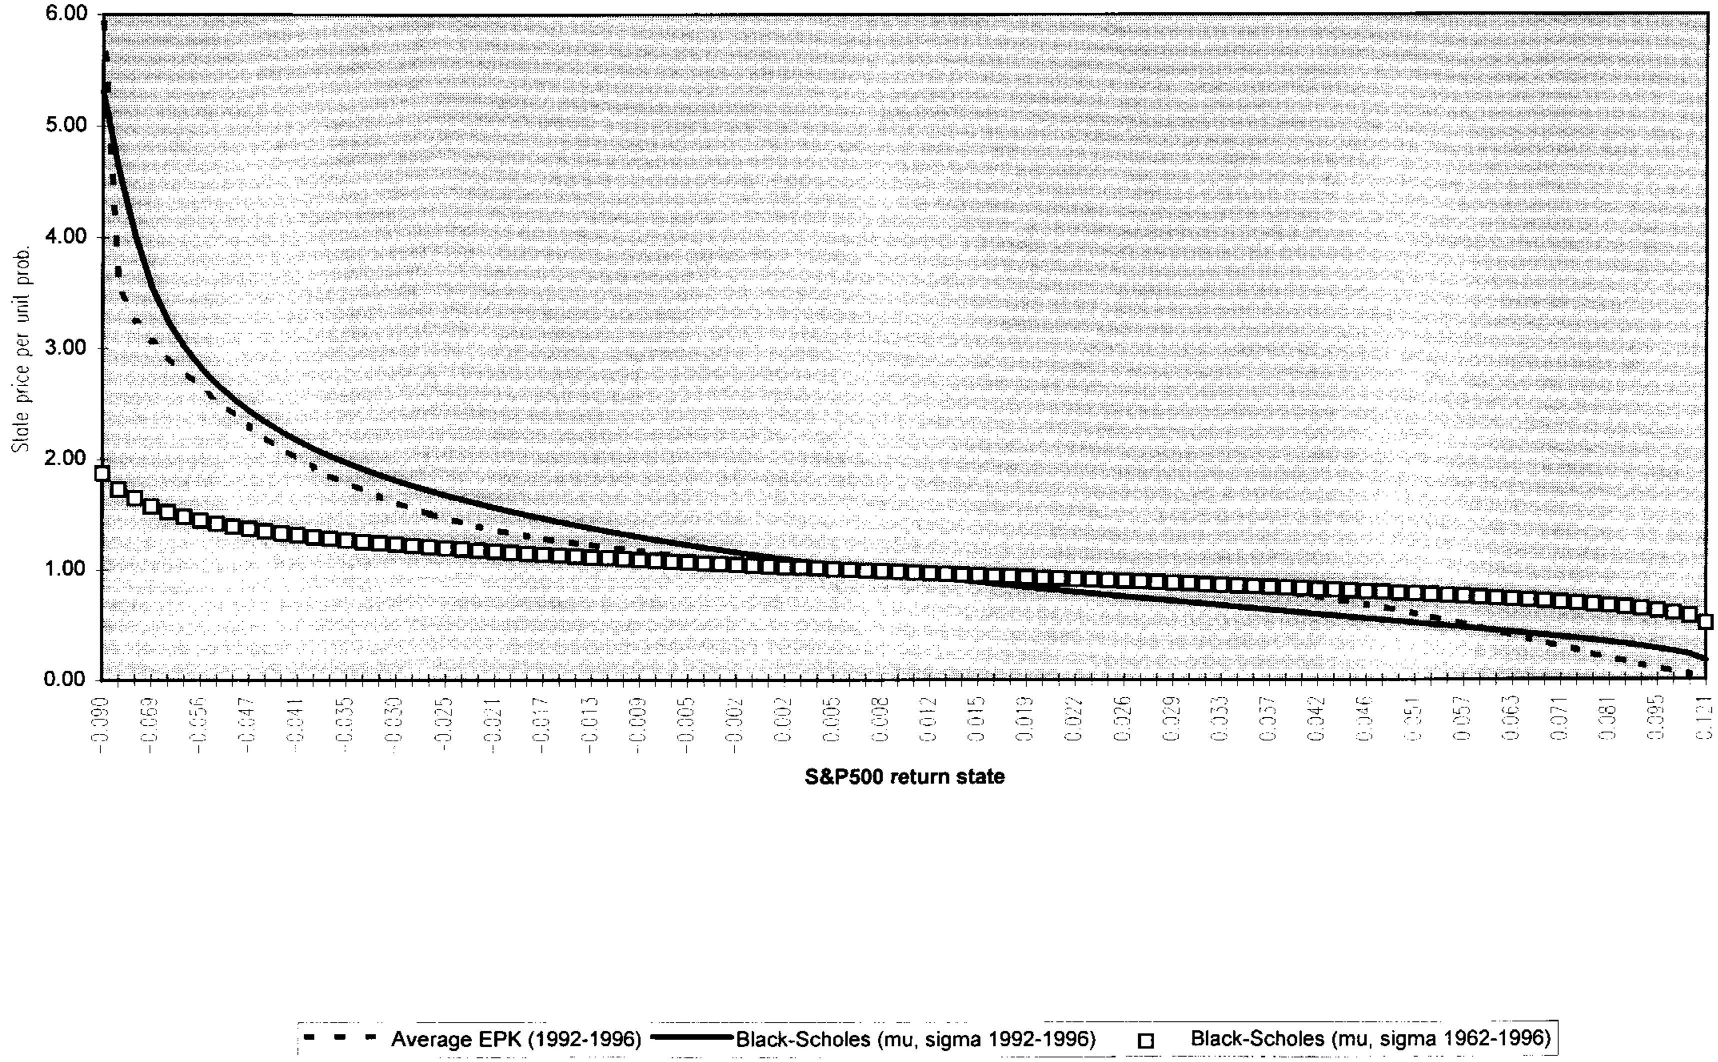

Figure 5
Daily empirical pricing kernels
May-June 1996

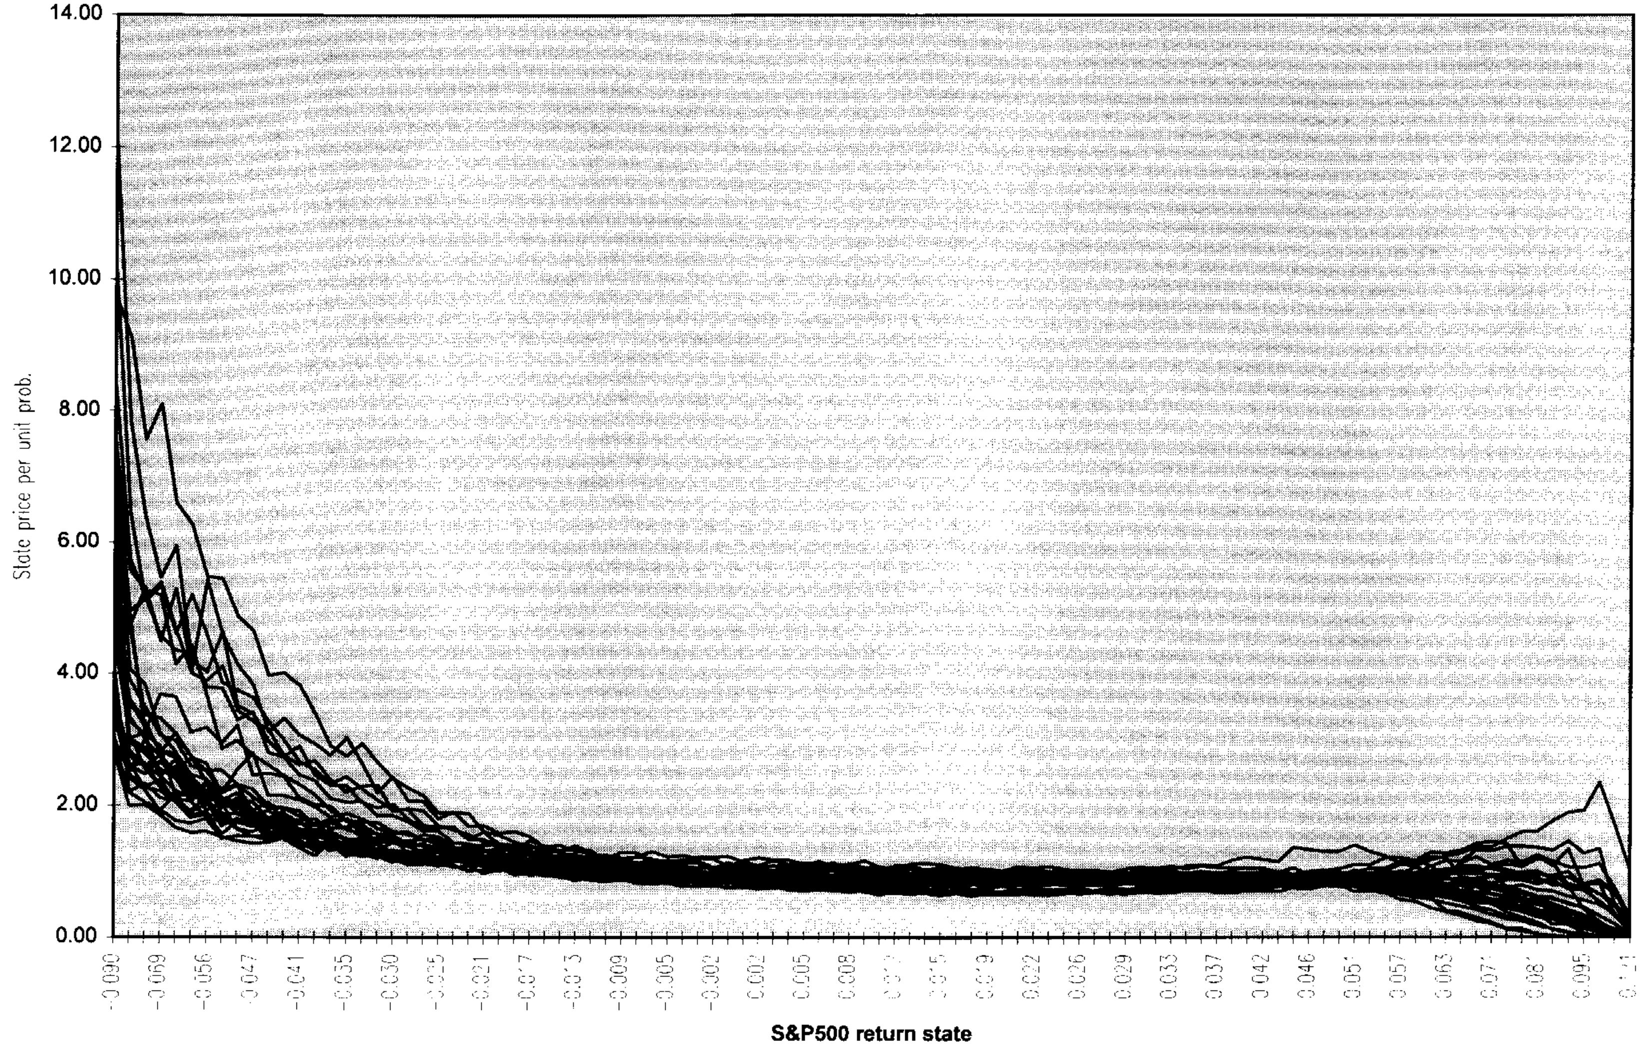

Figure 6
Average annual empirical pricing kernels, 1992-1996

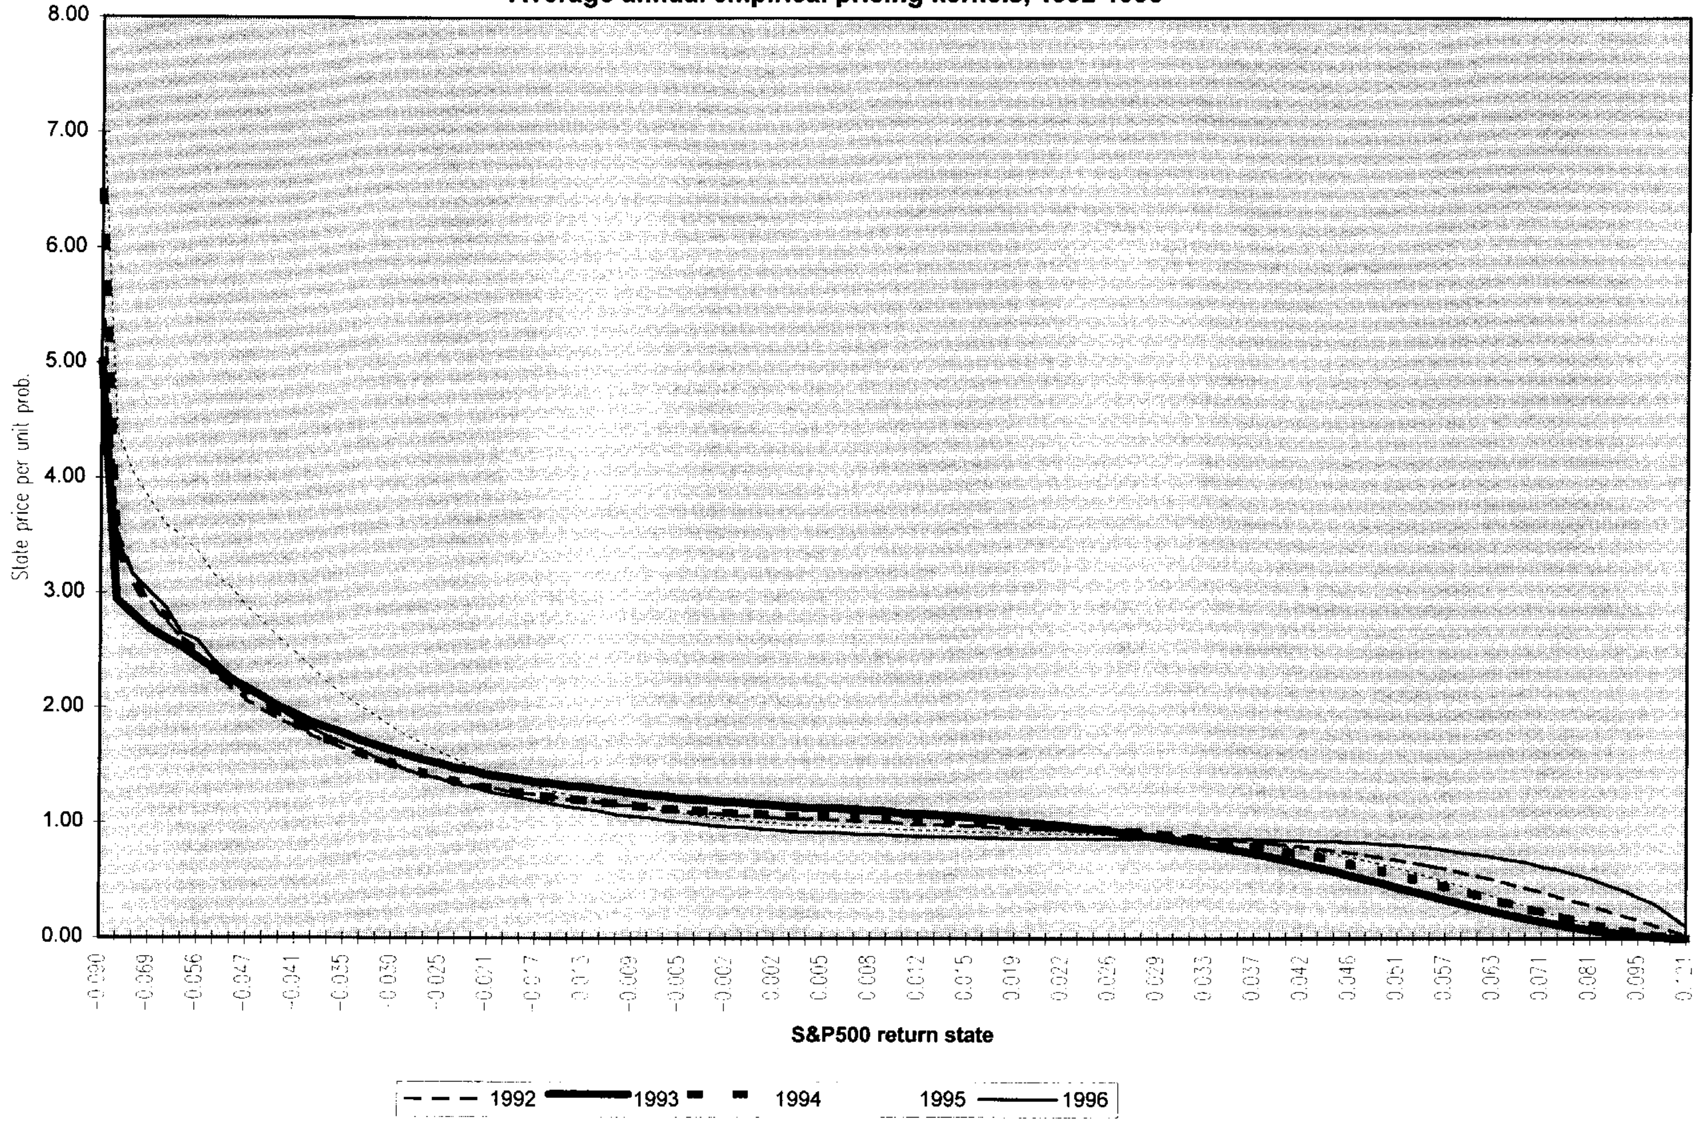

Figure 7
Power utility fit to average empirical pricing kernel
1992-1996

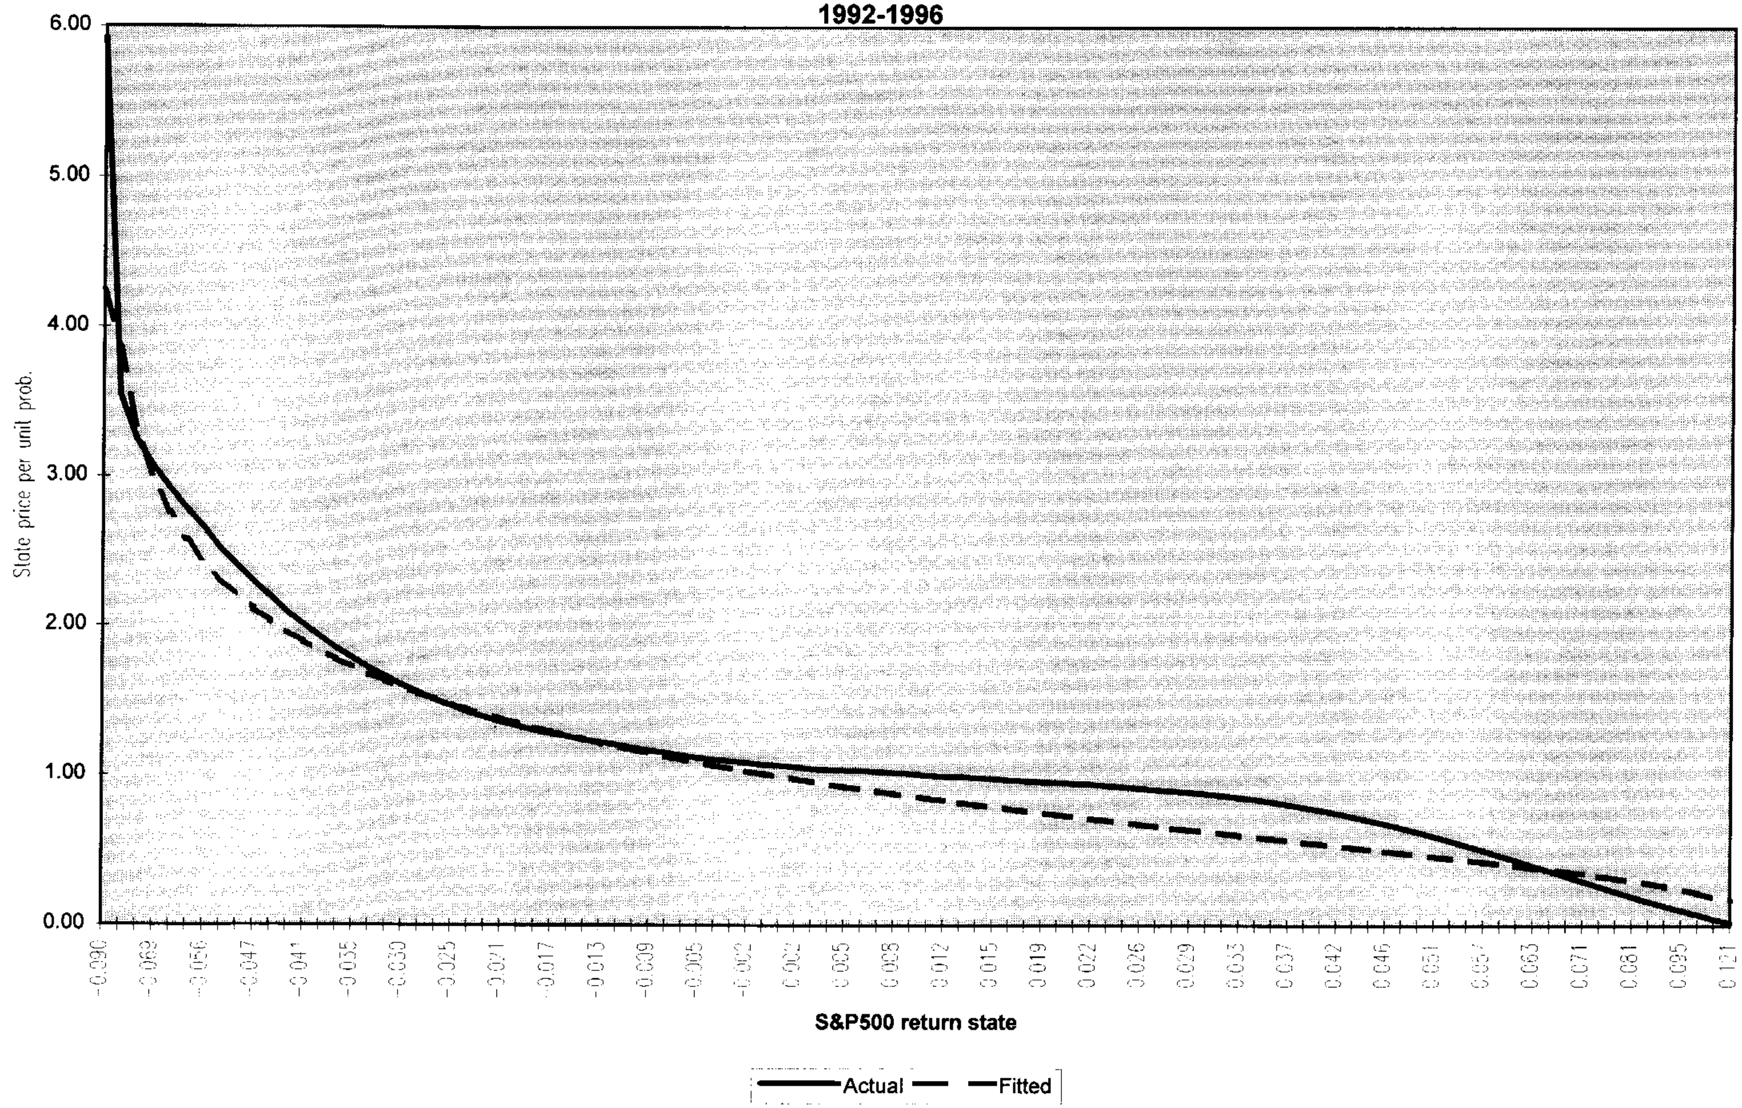

Figure 8
Estimated power utility gamma
Empirical pricing kernels, 1992-1996

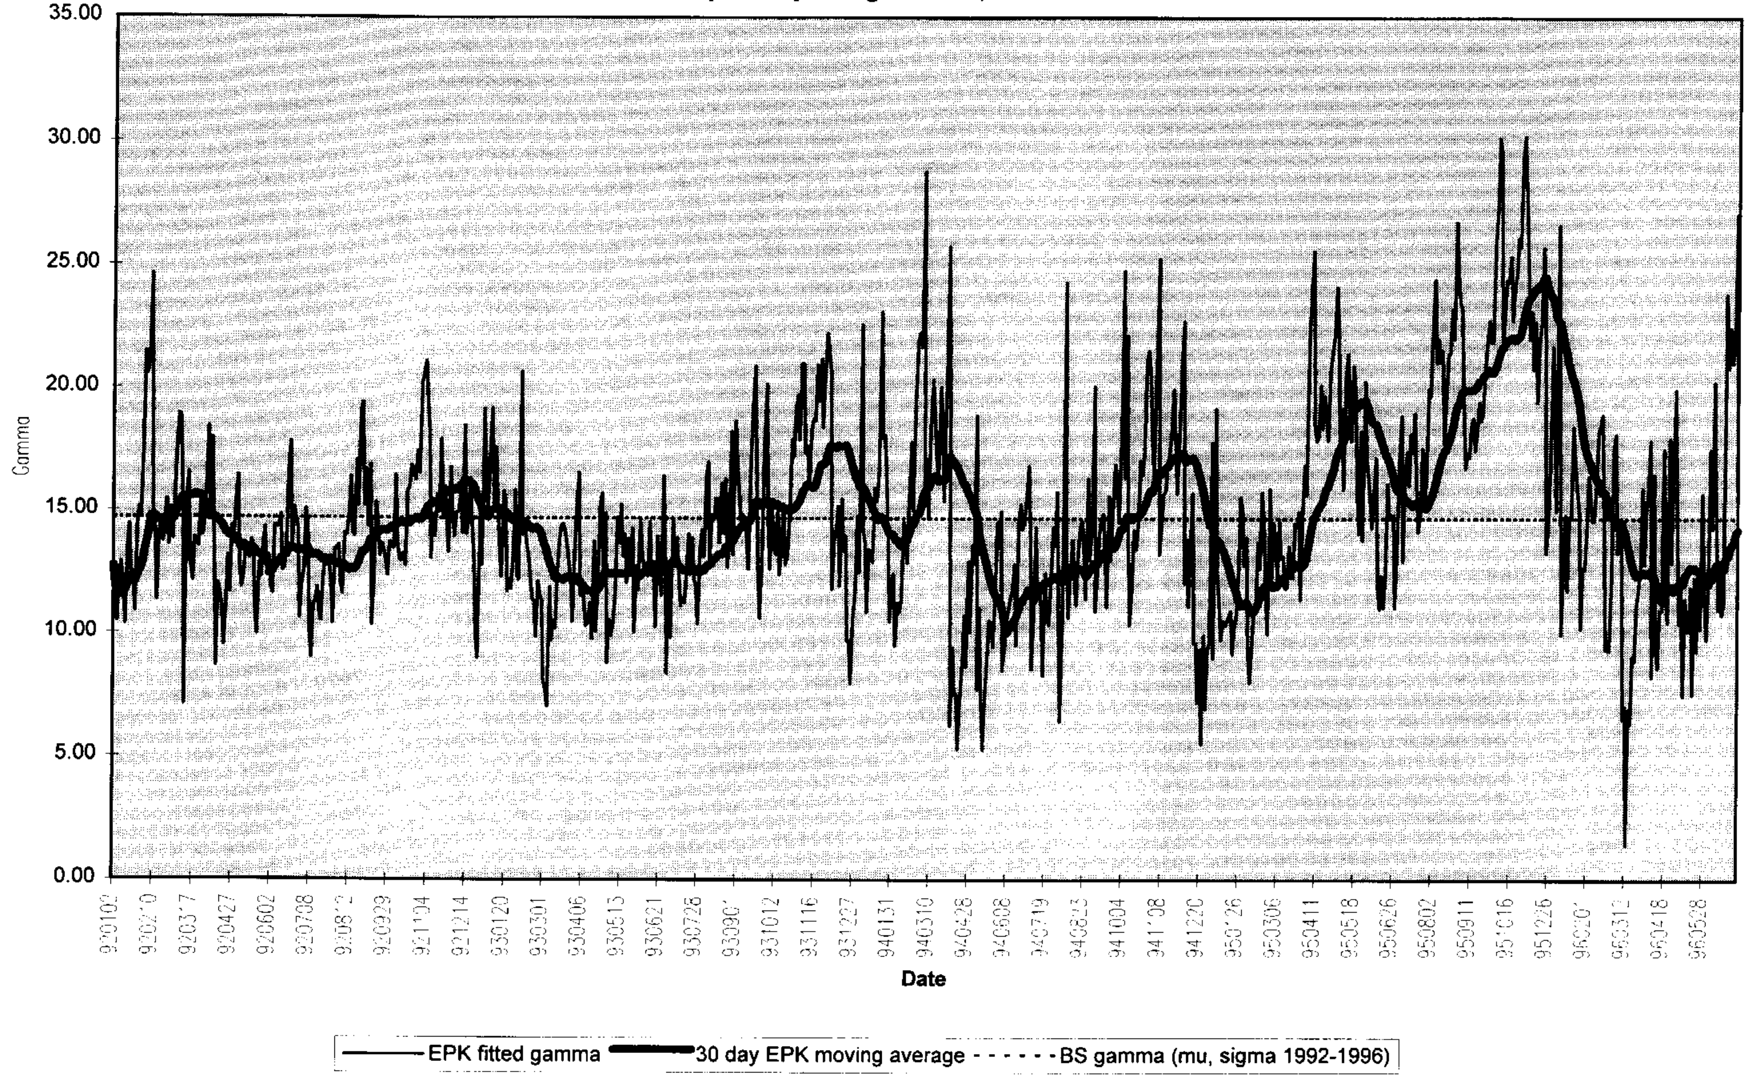

Figure 9
Preferences over crash state
T-t day return < -9%, 1992-1996