

Empirical pricing kernels

March 1999

Joshua V. Rosenberg
Department of Finance
NYU - Stern School of Business
44 West 4th Street, Suite 9-190
New York, New York 10012-1126
(212) 998-0311

Robert F. Engle
Department of Economics
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093-0508
(619) 534-1054

Abstract

This paper investigates the empirical characteristics of investor risk aversion over equity return states by estimating a daily semi-parametric pricing kernel. The two key features of this estimator are: (1) the functional form of the pricing kernel is estimated semi-parametrically, instead of being prespecified and (2) the pricing kernel is re-estimated on a daily basis, allowing measurement of time-variation in risk-aversion over equity return states.

Important empirical findings of the paper are as follows. Constant relative risk aversion over S&P500 return states is rejected in favor of a model in which relative risk aversion is stochastic. Empirical relative risk aversion over equity return states is found to be positively autocorrelated and positively correlated with the spread between implied and objective volatilities. In addition, the constant relative risk aversion (power utility) pricing kernel is found to underestimate the value of payoffs in large negative return states.

An option hedging methodology is developed as a test of the predictive information in the empirical pricing kernel and its associated state probability model. The results of hedging performance tests for out-of-the-money S&P500 index put options indicate that time-varying risk aversion over equity return states is an important factor affecting option prices.

This paper has benefited from the suggestions of Allan Timmerman, Bruce Lehmann, Steve Figlewski, Joel Hasbrouck, Ravi Bansal, Gur Huberman, Peter Carr, Jose Lopez, and participants in the 1998 Western Finance Association meetings, the 1998 Computational and Quantitative Finance conference, the NYU - Stern Finance department seminar, Columbia University Finance department seminar, the Federal Reserve Bank of New York research department seminar, and the 1997 Time Series Analysis of High Frequency Data Conference, an anonymous referee at the Journal of Financial Economics, and William Schwert. The original title for this paper (first draft: April 1997) was *Option hedging using empirical pricing kernels*.

I. Introduction

Payoffs, probabilities, and preferences are central elements of asset pricing. State contingent payoffs are typically predetermined by the contractual specifications of the asset, and state probabilities are estimated using data from the underlying price process. The least attention has been focused on estimating preferences, although misspecification of preferences may result in pricing and hedging errors.

The first-order condition associated with solution of the optimal consumption problem delineates the relationship between payoffs, preferences, probabilities, and prices. One of the most commonly utilized preference specifications is the power utility function defined over consumption states with coefficient of relative risk aversion equal to γ and rate of time preference equal to ρ . This implies the following pricing equation:

$$(1) \quad D_t = E_t \left[e^{-\rho(T-t)} (C_T / C_t)^{-\gamma} g(C_T) \right]$$

In this case, D_t is the current asset price denominated in units of consumption, C_t is the current level of consumption, C_T is the unknown future level of optimal consumption, and $g(C_T)$ is the payoff function which defines the number of units of consumption generated by ownership of the asset depending on the unknown future consumption state (C_T).

Equation (1) states that the asset price is the weighted expected payoff of the asset, where the conditional expectation is denoted by $E_t[\bullet]$. The weights are determined by the desirability of a unit payoff in each future consumption return state (C_T/C_t). Higher levels of relative risk aversion lead to greater value placed on payoffs in low consumption return states versus high consumption return states. In the asset pricing literature, the weighting function ($e^{-\rho(T-t)}(C_T/C_t)^{-\gamma}$) is often referred to as a pricing kernel.

While equation (1) is typically used to find the equilibrium asset price given payoffs, probabilities, and preferences; it is also natural to consider equation (1) as providing an implicit definition of preferences given observed payoffs, and probabilities, and observed market prices. Hansen and Singleton (1982, 1983) estimate the power pricing kernel parameters (ρ and γ) using a version of equation (1) in which both sides are divided by the current asset price and the conditional expectation is replaced with an unconditional expectation. An unconditional version of equation (1) is:

$$(2) \quad 0 = E\left[e^{-r(T-t)}(C_T / C_t)^{-\beta}(D_T / D_t) - 1\right]$$

Estimation is accomplished using the sample analog to equation (2), in which the sample time-series average replaces the unconditional expectation operator $E[\bullet]$. The representative investor's consumption return is estimated using per-capita real seasonally-adjusted consumption over a quarter (or month) as measured in the National Income and Products accounts. The gross asset return (D_T/D_t) is typically taken to be the observed return to a diversified U.S. equity portfolio.

More general techniques which characterize the pricing kernel without prespecification of its functional form have been proposed. Hansen and Jagannathan (1991) derive bounds for the mean and standard deviation of the consumption-based pricing kernel in terms of the mean and standard deviation of the market portfolio excess returns. Chapman (1997) extends the Hansen and Singleton technique to allow estimation of an approximate consumption-based pricing kernel pricing kernel of unknown functional form.

There are several potential weaknesses in these approaches. First, analyses dependent on aggregate consumption data are affected by measurement problems. Ermini (1989), Wilcox (1992), and Slesnick (1998) discuss issues such as coding errors, definitional problems, imputation procedures, and sampling error. Ferson and Harvey (1992) consider problems introduced by the Commerce Department's seasonal adjustment technique; Breeden, Gibbons, and Litzenberger (1989) address problems induced by use of time-aggregated rather than instantaneous consumption.

Second, analyses dependent on aggregate consumption data are not suited to identification of time-variation in pricing kernel parameters, since consumption is observed (at best) at a monthly frequency requiring a long consumption time-series for precise parameter estimation. Time-variation in pricing kernel parameters is useful to measure as an indicator of misspecification of the pricing kernel functional form or existence of omitted pricing kernel state variables.

Several papers have avoided this problem by substituting consumption data out of the pricing kernel in place of market variables. Brown and Gibbons (1985) replace the consumption return with a proxy for the market portfolio return under the assumption that that consumption is a fixed fraction of wealth. Bansal and Viswanathan (1993) assume that the market return and an interest rate are pricing kernel state variables instead of the measured consumption return. In their paper, semi-nonparametric estimation is used along with a version of equation (2) to identify the pricing kernel functional form.

The papers cited above do not make efficient use of asset price data. In particular, the unconditional version of equation (1) — represented by equation (2) — is used as an identifying condition. Equation (1) provides an exact pricing relationship, based on a conditional expectation, which may be used to more precisely estimate preference parameters. For example, Gallant, Hansen, and Tauchen (1990) show that the conditional mean and standard deviation of excess market returns may be used to improve upon the Hansen and Jagannathan (1991) pricing kernel bounds.

Recently, two papers — Jackwerth (1997) and Ait-Sahalia and Lo (1997) — have used a version of equation (1) to non-parametrically estimate the functional form of the pricing kernel as defined over equity return states. Their estimated pricing kernels may be viewed as a projections of the true pricing kernel — which might be defined over aggregate consumption— onto equity return states (S_T/S_t).

Since equity index options have payoff functions which depend explicitly on the future equity index level (S_T), they are useful in identifying a return-based pricing kernel. The following special case of equation (1) applied to the valuation of an equity index option reveals how the CRRA pricing kernel parameters are implicitly defined by index option prices.

$$(3) \quad D_t = E_t \left[e^{-r(T-t)} (S_T / S_t)^{-g} g(S_T) \right]$$

There are several reasons why a pricing kernel defined over equity return states is of interest. First, if consumption is a fixed fraction of wealth as represented by the level of an equity index, the return-based pricing kernel will approximately equal the consumption-based pricing kernel. In a single period model, the pricing kernels will be identical. Second, it can be shown (see appendix A), that if the equity index process is a geometric brownian motion (e.g. Black and Scholes, 1973), then the pricing kernel state variable is the equity index gross return. Finally, if the market return is a better proxy of the true consumption return than the data from the National Income and Product Accounts, pricing kernel estimation over equity return states may yield more insight into investor risk-aversion than estimation over NIPA consumption states.

Both Jackwerth (1997) and Ait-Sahalia and Lo (1997) estimate a *time-averaged* pricing kernel over equity return states. Average state prices are estimated using a pooled cross-section and time-series of option premia, and a smoothed historical returns density is used to estimate the unconditional return state probability density. The ratio of the average state prices and average state probabilities is

the average pricing kernel. Normalizing each average state price by its corresponding average state probability allows the value of each state contingent payoffs to be properly measured.

These papers extend the work of Rubinstein (1994) and Ait-Sahalia and Lo (1998) in which the state price density is found to exhibit significant deviations from log-normality. Since the state price density embeds information about state prices and state probabilities, it is unclear from this work whether this result is due to deviations from lognormal probabilities, deviations from a CRRA pricing kernel, or both factors.

A weakness in both of these papers is that neither uses the option data as efficiently as possible. A cost of the non-parametric estimation techniques which are used is that a large dataset of option premia is required to estimate a single pricing kernel. For this reason, option data is aggregated over time and the resulting average state prices are paired with unconditional state probabilities, to obtain an average pricing kernel.

With an average pricing kernel, it is not possible to detect variation in risk-aversion over *equity return states* which would occur in the case of decreasing relative risk aversion or habit formation (Campbell, 1996). In fact, it is only in the special case of a constant relative risk aversion pricing kernel over equity return states that empirical measurements of the state price per unit probability over these states will be constant over time. Jackwerth allows some time variation in state prices over several year periods without a corresponding model of time-varying state probabilities. This results in negative risk aversion estimates. With an average pricing kernel, it is also not possible measure economic factors which are correlated with the changes in level of risk-aversion.

The applicability of an average pricing kernels for pricing and hedging is limited, since assets will be correctly priced only when risk aversion and state probabilities are at their long-run average level. When either deviates its average, inaccurate prices (and hedge ratios) will result. There is substantial evidence of time-varying equity return state probabilities; see, for example, the literature on equity index return stochastic volatility summarized in section II.c.

Evidence for time variation in the equity return state price density has also been documented. For example, Jackwerth (1997) and Rubinstein (1998) present evidence that the estimated S&P500 state price density changed dramatically after the October 1987 stock market crash. Ait-Sahalia and Lo (p. 16) also acknowledge the problem of time-variation in the state price density (SPD): “In contrast, the kernel SPD estimator is consistent across time [emphasis added] but there may be some dates for which the SPD estimator fits the cross section of option prices poorly and other dates for which the SPD estimator performs very well (p. 16).”

This paper investigates the empirical characteristics of investor risk aversion over equity return states by estimating a daily semi-parametric pricing kernel. The two key features of this estimator are: (1) the functional form of the pricing kernel is estimated semi-parametrically, instead of being prespecified and (2) the pricing kernel is re-estimated on a daily basis, allowing measurement of time-variation in risk-aversion over equity return states.

The daily empirical pricing kernel (EPK) — which measures the state price per unit probability — is estimated as the ratio of state prices and state probabilities. In contrast to Jackwerth (1997) and Ait-Sahalia and Lo (1997) which estimate a time-averaged state price density, this paper estimates a daily state price density using the same-day cross-section of option prices. Thus, the EPK is compatible with a time-varying equity return state price density.

Also, in contrast to Jackwerth (1997) and Ait-Sahalia and Lo (1997) which estimate average state probabilities, this paper estimates state probabilities using current information and a model of the equity return process which incorporates stochastic volatility and non-normal innovations (e.g. jumps). Thus, the EPK is consistent with empirical evidence that equity return state probabilities (e.g. risk) are time-varying.

Important empirical findings of the paper are as follows. Constant relative risk aversion over S&P500 return states is rejected in favor of a model in which relative risk aversion is stochastic. Empirical relative risk aversion over equity return states is found to be positively and positively correlated with the spread between implied and objective volatilities. In addition, the constant relative risk aversion (power utility) pricing kernel underestimates the value of payoffs in large negative return states.

An option hedging methodology is developed as a test of the predictive information in the empirical pricing kernel and its associated state probability model. The EPK hedge is compared with a hedge based on constant relative risk aversion pricing kernel and lognormal state probabilities and a hedge based on the average empirical pricing kernel and average state probabilities. The second hedge uses state prices and state probabilities analogous to Jackwerth (1997) and Ait-Sahalia and Lo (1997).

In hedging tests of out-of-the-money S&P500 index put options, the EPK hedge outperforms the CRRRA hedge indicating that the power utility functional form is misspecified. The EPK hedge also outperforms the average EPK hedge affirming the presence of time-variation in the pricing kernel.

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

Estimation the empirical pricing kernel is based on the solution of the inverse of the asset pricing equation. In particular, the goal is to find a pricing kernel which rationalizes and observed set of asset prices, given the contractually specified state-contingent payoffs and set of estimated state probabilities.

Consider the following generalized version of equation (3), in which $K_{t,T}(S_T/S_t)$ represents the pricing kernel defined over equity return states, $g(S_T)$ is the payoff function for a particular security, and D_t is the price of the security.

$$(4) \quad D_t = E_t[K_{t,T}(S_T / S_t)g(S_T)]$$

Suppose the securities used in the estimation of the EPK are equity index options with expiration date T , and the equity index level on date T is S_T . Then, these derivative prices, along with equation (4), provide a rich set of identifying conditions for the pricing kernel.

Many pricing kernel estimation strategies are possible at this point. For example, in Rosenberg (1998), the parameters of a parametric pricing kernel are estimated by numerical inversion of equation (4). The approach taken in this paper is more general in that no particular function form is assumed for the pricing kernel.

In this paper, a semi-parametric estimation approach is utilized which is based on a discretization of the state space and separate estimation of state prices and probabilities. Since the pricing kernel is interpreted as the state price per unit probability, the ratio of state prices and state probabilities is used to estimate the pricing kernel. Discretization allows estimation of the EPK without postulating a specific pricing kernel functional form.

Let s_i be the subset of the real line corresponding to the return range for state i . With the state space divided into I states, the EPK is the ratio of the estimated discrete state price for state i , $z_{t,T}(s_i)$, with the estimated discrete probability of state i , $p_{t,T}(s_i)$. The discrete state prices may be interpreted as the prices of the supershares of Hakansson (1976) or the delta securities of Breeden and Litzenberger (1978).

$$(5) \quad K_{t,T}(s_i) = \frac{z_{t,T}(s_i)}{p_{t,T}(s_i)} \quad i = 1 \dots I$$

Notice that the pricing kernel is potentially time-varying over the discrete states, and it may implicitly depend on many economic variables such as current and lagged aggregate consumption, the level of interest rates, and so forth. Although equation (5) provides a non-parametric estimate of the EPK in the sense that the functional form of the EPK is unrestricted, restrictions on state prices and probabilities are reflected in the EPK, giving the estimation technique a semi-parametric interpretation.

II.b. Estimating state prices

State prices measure the value to the representative investor of an Arrow-Debreu security that pays one dollar in one state of the world and nothing in all other states of the world. As such, state prices reflect beliefs about the probability of the state occurring and the desirability of a payoff in the particular state of the world (i.e. risk-aversion). The first step in measurement of the empirical pricing is estimation of the discrete state prices, $z_{t,T}(s_i)$.

In this paper, states of the world are defined by the returns of an equity index portfolio. This particular state definition — also used in Jackwerth (1997) and Ait-Sahalia and Lo (1997) — is chosen so that asset payoff states are identical to pricing kernel states. This is convenient for estimation as is illustrated in equation (4).

If the pricing kernel were defined over a different set of states than the payoff variable, a mapping between the two state variables would be required to utilize the pricing equation. Also, as mentioned in the introduction, a returns-based pricing kernel may be preferred to the consumption-based pricing kernel as a measure of investor risk aversion, since aggregate consumption data is subject to substantial measurement error.

In this paper, a semi-parametric state price density estimation technique is utilized, which is based on the relationship between the state price density and derivatives of the call option pricing formula with respect to the option exercise price. Breeden and Litzenberger (1978) show that the first derivative is related to the state price distribution function and the second derivative is the state price density function. The second relationship has been utilized for state price estimation by Shimko (1993) and Ait-Sahalia and Lo (1998).

Ait-Sahalia and Lo (1998) and Shimko (1993) propose a generalized semi-parametric call option pricing formula which is defined by the Black-Scholes (1973) formula with the volatility parameter replaced by an implied volatility function which depends on the option exercise price.

$$(6) \quad C_t = BS(S_t, K, T - t, r_f, \sigma_t(K))$$

Estimation approaches differ in the methodology used to estimate the implied volatility function, $\sigma_t(K)$. Ait-Sahalia and Lo (1998) use non-parametric kernel regression with a time-invariant implied volatility function while Shimko (1993) uses a polynomial regression of implied volatilities on exercise prices. The non-parametric methodology of Ait-Sahalia and Lo (1998) is attractive in that it is the least restrictive in terms of assumptions made about the functional form of the state price density, but it is also the most data intensive. Ait-Sahalia and Lo (1998) assume that the implied volatility function is constant over a one-year period so that the option data may be time-aggregated for estimation.

In this paper, the state price density is estimated each day using only the current cross-section of option prices. Thus, the estimated state price density reflects investors current state probability forecasts and current level of relative risk aversion rather than a long-run average. Since there are typically about fifteen observations per day, a semi-parametric estimation technique is used which provides some structure for the data which diminishes as the amount of data increases.

For this purpose, the implied volatility function is estimated by fitting a minimum-distance B-spline — de Boor (1978) — to the observed implied volatilities as a function of the option exercise price. The number of spline coefficients and order of the spline are determined as an increasing function of the number of observed data points. Both are set to one-half the number of observations. Implied volatilities at non-traded exercise prices are interpolated using a cubic polynomial least-squares fit of implied variances on exercise prices.

Shimko (1993) derives the relationship between the state price distribution function, i.e. the integral of the state price density function, and the first derivative of the generalized call pricing formula. Let $n(\bullet)$ be the standard normal density function, let $N(\bullet)$ be the cumulative standard normal density function, and let $d_2 = [\ln(S_t/K) + 1/2(r_f - \sigma^2(T-t))] / \sigma\sqrt{T-t}$.

$$(7) \quad e^{-r_f(T-t)} \text{Prob}^*(S_T \leq K) = e^{-r_f(T-t)} [1 + K * n(d_2) * \sigma_{K,t}(K) - N(d_2)]$$

In equation (7), $\text{Prob}^*(S_T \leq K)$ is the risk-neutral distribution function and $e^{-r_f(T-t)}\text{Prob}^*(S_T \leq K)$ is the state price distribution function. The unknown function in this equation is $\sigma_{K,t}(K)$ — the first derivative of the volatility function with respect to the exercise price — which is estimated in this paper by numerically evaluating the first derivative of the spline interpolant with respect to the exercise price. Shimko (1993) uses a polynomial volatility function so that the derivative is obtained analytically.

Given the 100 discrete return states used in the empirical work, the discrete state prices for states 2 through 99 are obtained by evaluating equation (7) at the upper bound ($s_{u,i}$) and lower bound ($s_{b,i}$) of the return state (s_i) and taking the difference between the two values. States 1 and 100 are estimated using the cumulative probability below the lower bound of state 2 and the cumulative probability above the upper bound of state 99.

$$(8) \quad \begin{aligned} z_{t,T}(s_i) &= e^{-r_f(T-t)} \left[\text{Prob}^*(S_T \leq s_{l,2}) \right] & i = 1 \\ z_{t,T}(s_i) &= e^{-r_f(T-t)} \left[\text{Prob}^*(S_T \leq s_{u,i}) - \text{Prob}^*(S_T \leq s_{l,i}) \right] & i = 2 \dots 99 \\ z_{t,T}(s_i) &= e^{-r_f(T-t)} \left[\text{Prob}^*(S_T \geq s_{u,99}) \right] & i = 100 \end{aligned}$$

II.c. Estimating objective state probabilities

Three of the important features of equity index return process, which have been found in numerous studies (see, e.g. Ghysels, Harvey, and Renault, 1996), are that return volatility is stochastic and mean-reverting, return volatility increases disproportionately when the lagged return is negative, and return innovations are non-normal (e.g. there are return jumps). A stochastic volatility model of equity index return state probabilities should incorporate these features of the data.

In a discrete-time setting, stochastic volatility is most often modeled using extensions of the autoregressive conditional heteroskedasticity (ARCH) model proposed by Engle (1982). Comprehensive surveys of ARCH and related models are given by Bollerslev, Chou, and Kroner (1992) as well as Bollerslev, Engle, and Nelson (1994). In a continuous-time setting, stochastic volatility diffusions are commonly used. Surveys of this literature include Ghysels, Harvey, and Renault (1996) and Shephard (1996).

Comparative studies find that the ARCH and stochastic volatility diffusions perform similarly for many datasets. For example, Taylor (1994) as well as Kim, Shephard, and Chib (1998) find that ARCH and discretized diffusion models perform similarly in forecasting exchange rate volatility. This is not a surprising result since Nelson (1990) derives stochastic volatility diffusion limits for some classes of ARCH models. In empirical work, ARCH models are often preferred to SV diffusions because of their relative simplicity in estimation and their adaptability.

In ARCH (and related) models applied to financial returns, the conditional return variance is written as a function of lagged returns. In other words, these models postulate that current return variance depends on the magnitude, and sometimes the level, of recent returns. Since ARCH models are written in terms of the conditional variance, the contemporaneous error is implicit. However, both volatility forecasts and volatility realizations are stochastic, since both depend on realized squared error terms which are stochastic. Baillie and Bollerslev (1992) analyze the prediction error distributions in generalized ARCH (GARCH) models.

The adaptability of ARCH framework is illustrated by the enormous class of extended ARCH models which have been developed to incorporate relevant features of volatility dynamics in different time series. For example, Bollerslev (1986) generalizes the ARCH structure and proposes a GARCH model in which the current volatility is predicted by its own lag as well as the lagged squared innovation. The relationship between an ARCH(1) and GARCH(1,1) model is analogous to that of an MA(1) and an ARMA(1,1) model. Engle and Lee (1993) introduce a restricted version of a GARCH(2,2) model which they refer to as the GARCH components model. In this model, volatility has a short-run and long-run factor, both of which are mean reverting but at different rates.

The asymmetric feature of equity return volatility, first noted by Black (1976), has been incorporated into the ARCH framework. See, for example, Glosten, Jagannathan, and Runkle (1993) and Zakoian (1994). ARCH models have also been generalized to allow for non-normal return innovations. Bollerslev (1987) proposes a GARCH model with t-distributed innovations, while Engle and Gonzales-Rivera (1991) derive a semi-parametric model of the innovation density.

An alternative approach for estimating state probabilities — smoothing the historical return histogram — is used Jackwerth (1997) and Ait-Sahalia and Lo (1997). However, their techniques are only appropriate for estimating average state probabilities rather than time-varying state probabilities. By construction, their methods result in state probabilities that are constant over the estimation period.

For example, over the estimation period used in these studies, S&P500 return variance is held constant. Thus, options are priced each day using the same volatility for a period of months or years.

Neither of these techniques is appropriate for the study of daily variation in risk aversion. If state probabilities are assumed to be constant, then observed daily changes in option prices will be attributed to daily changes in risk-aversion. In fact, daily changes in option prices may solely be due to changes in equity index volatility. This is the reason that a dynamic state probability model, which captures time-variation in state probabilities, is used in this paper.

The three important features of equity index return process are modeled in this paper using an asymmetric GARCH model with an estimated innovation density. This stochastic volatility model incorporates asymmetric volatility effects using a specification based on Glosten, Jagannathan, and Runkle (1993). This model also utilizes an estimated innovation density which captures the potential non-normalities of the true innovation density. The correctness of this model specification for S&P500 returns is analyzed in section III.a.

The asymmetric GARCH model is specified as follows:

$$(9) \quad \ln(S_t / S_{t-1}) = \mathbf{m} + \mathbf{e}_t, \quad \mathbf{e}_t \sim D_{empirical}(0, \mathbf{s}_{t|t-1}^2)$$

$$(10) \quad \mathbf{s}_{t|t-1}^2 = \mathbf{w} + \mathbf{a}\mathbf{e}_{t-1}^2 + \mathbf{b}\mathbf{s}_{t-1|t-2}^2 + \mathbf{g}\mathbf{Max}[0, -\mathbf{e}_{t-1}]^2$$

Equation (9) states that the daily log-return ($\ln(S_t/S_{t-1})$) has a constant mean (μ) and innovations (ϵ_t) drawn from an empirical density function ($D_{empirical}$) which has time-varying stochastic variance ($\sigma_{t|t-1}^2$). While a constant expected return is not usually compatible with time-varying state prices, the time-variation in expected returns over the short time horizon (about 20 days) considered in this paper is not likely to be important, and equation (9) may be viewed as an approximation. Equation (10) states that the daily conditional return variance ($\sigma_{t|t-1}^2$) depends on a constant (ω), the one-day lagged squared innovation ($-\epsilon_{t-1}^2$), and whether the lagged return was negative ($\mathbf{Max}[0, -\epsilon_{t-1}]$).

This stochastic volatility model is estimated by maximum likelihood, under the assumption that the error density is normal. Bollerslev and Wooldridge (1992) provide conditions under which this estimation technique has a quasi-maximum likelihood interpretation and generates consistent parameter estimates even when the true error density is non-normal. The estimated innovation density ($D_{empirical}$) is taken to be the density function defined by dividing the estimated return innovations by their estimated conditional variances. This empirical density will reflect excess skewness and kurtosis which is not explained by the stochastic volatility model.

For the asymmetric GARCH model, the conditional state probabilities may be obtained using a Monte-Carlo approach. A consistent estimate of probability of a given return state occurring is the proportion of the total number of simulated returns ($r_{t,T,j}$, $j=1 \dots N$) that fall within a state return range (s_i). In particular, let $p_{t,T}(s_i)$ be the estimated probability of state s_i occurring; and let $I(r_{t,T,j} \in s_i)$ be an indicator function that takes the value one when the simulated return falls within the state return range defined by s_i and zero otherwise.

$$(11) \quad p_{t,T}(s_i) = \frac{1}{N} \sum_{j=1}^N I(r_{t,T,j} \in s_i) \quad I(\bullet) \text{ is an indicator function} \quad i = 1 \dots I$$

In equation (11), the discrete state probabilities are obtained by simulating the estimated asymmetric GARCH process $N=200,000$ times to generate the $j=1 \dots 200,000$ returns ($r_{t,T,j}$), with innovations drawn from the estimated error density.

II.d. Estimating EPK hedge ratios

In this paper, we construct hedge ratios which are designed to neutralize the portfolio to first and second-order effects of changes in the underlying price. In a continuous-time diffusion setting, this type of hedge is sufficient to eliminate all randomness in the hedge portfolio and also provides a minimum-variance hedge.

When there are additional random factors such as stochastic volatility or stochastic relative risk aversion, a perfect hedge also requires that the option portfolio exposure to these factors be measured and hedged. The proposed first and second-order hedge ratios will hedge changes in volatility, if volatility is related (linearly or nonlinearly) to underlying price changes. However, the hedges will reduce, but not eliminate hedge portfolio variability.

Consider the following Taylor series expansion which measures the option price change due to first and second order effects underlying price changes.

$$(12) \quad D_{t+1} - D_t \cong \frac{\mathcal{D}_{t+1}}{\mathcal{S}_{t+1}} (S_{t+1} - S_t) + \frac{1}{2} \frac{\mathcal{D}^2_{t+1}}{\mathcal{S}_{t+1}^2} (S_{t+1} - S_t)^2$$

In equation (12), the first and second derivatives are measures of exposure to changes in the underlying price. For example, to eliminate exposure of a portfolio containing a written out-of-the-money put option to first-order underlying price effects, an appropriate number of units (the hedge ratio) of the underlying asset or at-the-money put should be purchased. The hedge ratio is given by the opposite of the ratio of the exposure of the instrument to be hedged and the exposure of the hedging instrument.

Unfortunately, the option pricing formula given in equation (4) does not facilitate calculation of option price derivatives in closed form. Bates (1995) proposes a method of extracting these derivatives which is consistent with an estimated implied volatility function. However, Bates' method is only correct in the case in which stock returns are uncorrelated with return volatility — an assumption which is violated in the case of equity index options.

Engle and Rosenberg (1995) introduce a technique for calculating approximate option price derivatives using monte-carlo simulation and centered finite differences in a setting with risk-neutral investors. This paper generalizes the Engle and Rosenberg technique to a setting with stochastic risk-aversion by incorporating the pricing kernel into the option pricing formula.

Following Engle and Rosenberg (1995), consider three possible one-day underlying price changes. The stock price could rise by one-standard deviation to $S_t + \varepsilon$, remain constant at S_t , or fall by one standard deviation to $S_t - \varepsilon$. Each change results in a different realization for the derivative price ($\hat{D}_{t+1|S_{t+\varepsilon}}, \hat{D}_{t+1|S_t}, \hat{D}_{t+1|S_{t-\varepsilon}}$). These derivative prices may be used to calculate centered finite difference approximations to the first and second derivatives of the option pricing formula as illustrated in equation (13).

$$(13) \quad \frac{\mathcal{D}_{t+1}}{\mathcal{S}_{t+1}} \cong \frac{\hat{D}_{t+1|S_{t+\varepsilon}} - \hat{D}_{t+1|S_{t-\varepsilon}}}{2\varepsilon} \quad \frac{\mathcal{D}^2_{t+1}}{\mathcal{S}_{t+1}^2} \cong \frac{\hat{D}_{t+1|S_{t+\varepsilon}} - 2\hat{D}_{t+1|S_t} + \hat{D}_{t+1|S_{t-\varepsilon}}}{\varepsilon^2}$$

The date t+1 derivative prices are evaluated using the sample analog of equation (4). The derivative price at date t+1, conditional on the underlying price at date t+1, is expectation of the pricing kernel weighted payoff evaluated by monte-carlo simulation.

$$(14) \quad \hat{D}_{t+1|S_{t+1}} = \frac{1}{N} \sum_{n=1}^N \sum_{i=1}^I K_{t+1,T}(s_i) g(s_i) I(r_{t+1,T,n} \in s_i)$$

In equation (14), the sample average replaces the integral used in equation (4). Preferences are represented by the one-step ahead forecast of the empirical pricing kernel $K_{t+1,T}(s_i)$, payoffs are represented by the payoff function $g(s_i)$, and probabilities are represented by simulated realizations from a dynamic probability model given the date $t+1$ underlying price (S_{t+1}) and the n^{th} simulated return state outcome ($r_{t+1,T,n}$).

In the case of the empirical pricing kernel, one-step ahead forecast for each state price per unit probability is used to estimate tomorrow's pricing kernel ($K_{t+1,T}$). The one-step ahead forecast is obtained using an autoregressive model with two lagged values estimated for each of the one-hundred state price per unit probabilities.

The first summation in equation (14) is over the N simulated return states defined by returns from date $t+1$ to date T . The second summation is over the I states of the world, and the function $I(r_{t+1,T,n} \in s_i)$ is an indicator function. When the n^{th} return falls within the i^{th} state return range, the indicator function is unity and the state return is weighted by the pricing kernel at the i^{th} state. For the n^{th} return in all other states, the indicator function is zero.

Equations (13) and (14) may also be used to estimate hedge ratios for a constant relative risk aversion pricing kernel - lognormal probability model and an average EPK - average state probability model with the appropriate substitutions of pricing kernels and state probability models.

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

S&P500 state prices and probabilities are estimated on 1235 of 1517 days over the period from January 1991 through December 1995. Dates not included in the sample are those on which there are no bid or ask quotes for option contracts with between 10 and 60 days until expiration.

On each date, the nearest maturity S&P500 index option series with between 10 and 60 days until expiration is chosen. One quote for each option series is obtained by averaging the midpoint of the last bid and ask quote of the day. The intraday Black-Scholes implied volatility is calculated using the synchronized S&P500 index price measured using the last transactions price for the corresponding S&P500 futures contract and the cost-of-carry relationship, the riskless rate measured using the three month treasury bill yield, the dividend yield measured using the implied dividend yield from the put-call parity formula, and the time until expiration measured using the number of trading days until expiration.

The first panel of Table 1 reports the properties of option contracts used in state price estimation. On a typical day, there are 15 different contracts with reported price quotes, not including those which are informationally equivalent based on put-call parity. The average daily moneyness range is 12% with the moneyness range defined as is the sum of the absolute values of moneyness of farthest OTM call and the farthest OTM put. (Moneyness is $S_t/E - 1$ for calls and $E/S_t - 1$ for puts, where S_t is the closing S&P500 index level and E is the option exercise price). There are usually more contracts with exercise prices below the spot price than above. The average time-until-expiration is about 20 days.

States of the world are defined by S&P500 index returns. State 1 includes returns less than -9.85% , state 2 includes returns from -9.85% to -9.70% , states 3 through 99 include adjacent return ranges of $.15\%$, and state 100 includes returns greater than or equal to 5.00% . The asymmetric definition of the return states is based on the larger amount of information about the state price density below the spot price due to more options trading with exercise prices below the contemporaneous spot price.

The state price estimation procedure described in section II.b. is implemented for each of the 1235 dates. Figure 1 plots the estimated state price densities for the first estimation date of each year (1991-1995) for states 2 through 99. The time-variation in state-price densities is apparent with significantly higher state prices for large negative return states in 1991 and 1993.

The state price densities also exhibit excess negative skewness and positive kurtosis relative to a lognormal density. This is consistent with the results of Rubinstein (1994) and Ait-Sahalia and Lo (1998). At this stage, it cannot be determined whether the skewness in the state price density is due to higher probabilities associated with low return states or a pricing kernel which places high value on positive payoffs in these states.

State probabilities are obtained by simulation of the S&P500 return process using the asymmetric GARCH model described in section II.c. Table 2 reports summary statistics for the daily S&P500 index returns series used in the model estimation for the period 1990-1995. Over this period, the average annualized S&P500 index return is 11.5% , and the annualized S&P500 return standard deviation is 9.7% . S&P500 returns exhibit negative skewness and positive kurtosis, which is consistent with a stochastic volatility model with asymmetric effects.

Table 2 also describes the dynamic state probability model used to estimate the state probabilities. Three nested GARCH models are estimated: ARCH(1), GARCH(1,1), and asymmetric GARCH(1,1). A likelihood ratio test is used to compare the statistical significance of the increase in likelihood for each model generalization.

The GARCH model is found to offer a statistically significant improvement over the ARCH model with a likelihood ratio test p-value less than .0001. In addition, the asymmetric GARCH model is found to offer a statistically significant improvement over the GARCH model with a likelihood ratio test p-value less than .0001. This is confirmed by the significant robust t-statistic for the volatility asymmetry parameter (γ) of 2.24.

Two specification tests are used to analyze the correctness of the asymmetric GARCH specification. The first test — Engle's (1982) ARCH LM test — measures the presence of stochastic volatility effects have not been entirely explained by the model. The asymmetric GARCH model passes this test with a p-value of .9200. The second test — Ljung-Box (1978) Q-statistic — measures unexplained serial correlation in the model standardized residuals. The asymmetric GARCH model passes this test with a p-value of .2152. The asymmetric GARCH model reported in Table 2 is chosen as the preferred model and is used for state probability estimation.

The standardized residuals from this model are used to define the empirical distribution ($D_{\text{empirical}}$) of the innovations for simulation of the state probabilities. The standardized residuals deviate significantly from normality with negative skewness and positive excess kurtosis, as reported in the third panel of Table 2. Non-normality does not affect the conditional volatility estimates, but it does affect the higher moments of forecast state probability distribution.

Figure 2 plots the estimated state probability densities for the first estimation date of each year (1991-1995) for states 2 through 99. The state probability densities may be compared with the state price densities plotted in Figure 1. The time-variation in state probability densities is apparent with significantly higher state probabilities for large negative return states in 1991 and 1993.

Figure 3 plots the conditional volatility forecasts using the asymmetric GARCH model. Over this period, estimated annualized S&P500 volatility ranges from 7.1% to 24.12% with a standard deviation of 2.9%.

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

The empirical pricing kernel is estimated using the daily ratio of the estimated state prices and state probabilities as described in Section III.a. Each day, the empirical pricing kernel provides an estimate of the probability normalized value to a representative investor of a payoff in each discrete return state.

Figure 4 plots the empirical pricing kernel for the first estimation date of each year (1991-1995) for states 2 through 99. The daily EPK shapes are similar, but the daily level of risk aversion varies

substantially. The state price per unit probability for large negative return states is especially volatile reflecting time-varying demand for insurance against a significant market decline. If S&P500 returns are a proxy for aggregate consumption, the negative slope of the empirical pricing kernels indicates that investors value unit payoffs in low consumption states more than in high consumption states. The time variation in the EPK suggests that constant risk aversion over equity return states is not supported by the data.

Notice that on the first day of 1991 and 1993, both the state prices and the state probabilities for large negative return states were high. Thus, the resulting state prices per unit probability were not especially high on either date. This illustrates the importance of matching state prices with state probabilities which reflect investors beliefs at the same point in time.

Figure 5 compares the time-averaged empirical pricing kernel with a constant relative risk aversion pricing kernel estimated over the same period. The average EPK is obtained by taking a state-by-state average of the state price per unit probability over the estimation period from 1991 through 1995. This average EPK measures average preferences over the estimation period which is a similar measure to that of Ait-Sahalia and Lo (1997) or Jackwerth (1997).

We find that functional form for the average EPK is similar in shape to that of Ait-Sahalia and Lo (1997). In both studies, the value of payoffs in S&P500 crash states is many times greater than the value of payoffs in small return states. And, the value of payoffs is declining as wealth (S&P500 returns or levels) increase.

Our results are substantially different than Jackwerth (1997). He finds that after 1987, investors exhibit negative risk aversion (are risk-loving) for S&P500 returns close to zero and that absolute risk aversion is increasing as wealth increases. In addition, to being inconsistent with standard economic assumptions, neither of these results is similar with the results of this paper or Ait-Sahalia and Lo (1997). Jackwerth's (1997) results may be due to pairing state prices with an incorrect state probability model.

The constant relative risk aversion pricing kernel is estimated using the relationship between lognormal state prices, lognormal state probabilities, and constant relative risk aversion. Equation A.4 in Appendix A provides the formulas used for estimation. The input variables —the riskless rate, expected return, and volatility — are estimated using averages over the period from 1991-1995.

There are prominent differences between the constant relative risk aversion pricing kernel and the average empirical pricing kernel as illustrated in Figure 5. Most notably, the constant relative risk aversion pricing kernel underestimates the value of payoffs in large negative return states. A similar

result is found in Bansal and Viswanathan (1993) and Ait-Sahalia and Lo (1997). The average probability standardized value to an investor of a \$1 payoff in state 1 (S&P500 return less than -9.85%) is about \$36.15/per unit probability using the average EPK and about \$2.80/per unit probability using the CRRA pricing kernel.

This paper is unique in that time-variation in risk aversion over equity return states is measured. To analyze the time-series properties of the empirical pricing kernel, a daily summary measure of EPK relative risk aversion ($\gamma_{EPK,t}$) is constructed. Specifically, $\gamma_{EPK,t}$ is set equal to the exponent of a power function obtained by minimizing the distance between the power function and the EPK.

Using the power function to approximate the EPK facilitates a direct comparison of empirical and constant relative risk aversion coefficients. Equation (15) provides the definition of $\gamma_{EPK,t}$:

$$(15) \quad \underset{\xi_{EPK,t}}{\text{Min}} \sum_{i=1}^{100} \left[K_{i,T}(s_i) - (1+r_i)^{-\xi_{EPK,t}} \right]^2$$

In equation (15), r_i is discrete state return, which is defined as the midpoint of the state return range for states 2 through 99, $r_1 = -10.00\%$, and $r_{100} = 5.00\%$. Using equation (15), a daily level of risk aversion is obtained from 1991 - 1995.

Table 3 provides summary statistics describing the behavior of empirical relative risk aversion. Over this period, the average $\gamma_{EPK,t}$ is 22.14 and the range is from 1.21 to 57.99. Using S&P500 options data for 1993, Ait-Sahalia and Lo (1997) obtain an average implied risk-aversion coefficient of 25.5. We find a similar result; the average $\gamma_{EPK,t}$ is 24.11 over 1993. Figure 6 graphs the risk aversion dynamics from 1991 through 1995 as measured by $\gamma_{EPK,t}$.

For comparison, the coefficient of relative risk aversion from the CRRA pricing kernel (γ_{CRRA}) is obtained using equation A.3 in Appendix A. Using sample averages over the period 1991-1995 for evaluation of the formula, γ_{CRRA} is found to be 8.84. Daily empirical relative risk aversion ($\gamma_{EPK,t}$) is usually higher but sometimes lower than the coefficient of relative risk aversion (γ_{CRRA}). However, average empirical risk aversion is substantially higher than the coefficient of relative risk aversion.

Persistence in EPK risk aversion is illustrated by a first-order autocorrelation of .80 and second order partial autocorrelation of .16. It appears that an autoregressive model with two lags adequately describes the dynamics of EPK risk aversion.

Time variation in risk aversion may be used to identify sources of changes in risk aversion. The third panel of Table 3 contains correlations of EPK risk aversion with several other variables. S&P500 returns are not found to be correlated with EPK risk aversion which is evidence against decreasing or increasing relative risk aversion. Interest rate changes are also uncorrelated with EPK risk aversion.

Another variable considered is the change in the spread between implied and the GARCH forecast volatility. This change is found to be positively correlated with the level of risk aversion. This result is intuitive, since the spread is a proxy for the risk-aversion based on the amount the market is pricing options above the objective volatility.

Evidence from the analysis of the empirical pricing kernel confirms that empirical relative risk aversion follows a mean-reverting stochastic process with moderate persistence. Two market variables are tested and found to be uncorrelated with EPK risk aversion.

IV. Hedging tests

The purpose of the option hedging tests used in this section is to identify whether the empirical pricing kernel provides an effective measure of the level and dynamics of investor risk aversion. Hedge ratios depend on risk aversion and state probability estimates, so the accuracy of each pricing kernel - probability model may be judged by its hedging performance. To this end, three pricing kernel and probability models are compared: constant relative risk aversion - lognormal state probability, average empirical pricing kernel - average state probability, empirical pricing kernel - asymmetric GARCH state probability.

The hedging tests are based on hedging a \$100 portfolio out-of-the-money (OTM) S&P500 index put options using at-the-money (ATM) put options and the index. The hedging objective is to minimize the standard deviation of one-day hedge portfolio price changes. This 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.

The second 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 maturity 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 -3.6%, and the average time-to-maturity is approximately 20 trading days.

Synchronized end-of-day option price quotes used in the hedging tests are obtained by repricing the options using the dividend-adjusted Black-Scholes formula evaluated at the intraday implied volatility and the end-of-day S&P500 index level.

The average ATM and OTM daily implied return volatilities of 15.3% and 12.2% illustrate the implied volatility skew which characterizes the sample period. OTM implied volatilities are consistently higher than ATM implied volatilities. The standard deviations of ATM and OTM implied volatilities illustrate that implied volatilities are not constant over the sample period.

For the CRRA - lognormal model, the hedge ratios are obtained using the estimated CRRA pricing kernel (section III.b.) and hedge ratio estimation technique described in II.d. State probabilities are generated using the discretization corresponding to a geometric brownian motion with drift and diffusion parameters set to their sample values over the estimation period.

The average empirical pricing kernel - average state probability model hedge ratios are obtained using the time-averaged EPK (section III.b.) and hedge ratio estimation technique described in section II.d. State probabilities are generated by simulating independent identically distributed returns from the empirical distribution of S&P500 returns over the period from 1990 - 1995.

The empirical pricing kernel - asymmetric GARCH state probability model hedge ratios are obtained using the estimated EPK (section III.b.) and hedge ratio estimation technique described in section II.d. State probabilities are generated by simulating using the estimated asymmetric GARCH model and the estimated innovation density.

First and second-order exposures to underlying price changes for each model are generated using the centered finite difference approximations described in section II.d. Each one-day ahead derivative price is obtained using 100,000 simulation replications. Hedge ratios are calculated using the standard technique.

Table 4 reports the hedging test results. The results of using the three 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 \$27.00 per day. The most effective hedge using the underlying asset alone is the daily EPK hedge which reduces the portfolio standard deviation by about 40% to \$16.83 per day. The least effective hedge is the CRRA hedge which reduces the standard deviation by about 20% to \$20.37 per day. The average EPK hedge performance is in between the other two methods.

A test of the statistical significance of the improvement in hedging performance using the EPK hedge is performed using a predictive accuracy statistic similar to that proposed by Diebold and

Mariano (1995). The magnitude of the hedging error of two models is compared using the t-statistic for the absolute error loss, which is defined as the difference between the absolute value of the alternative model hedging error and the absolute value of the EPK hedging error. Standard errors are calculated using the heteroskedasticity and autocorrelation consistent covariance matrix of Newey and West (1987). When the t-statistic is significant and positive, the daily EPK hedge performance is significantly better than the alternative model hedge performance.

The fourth column of Table 4 reports this measure. The t-statistic for the daily EPK model outperformance is 3.47 against the average EPK model, which is statistically significant at standard significance levels. The t-statistic for the daily EPK model outperformance is 6.80 against the CRRA model, which is also statistically significant.

The pattern of hedge performance is similar for hedges using the ATM put alone. The daily EPK hedge outperforms the average EPK and CRRA alternatives, and the differences are statistically significant. In addition, hedging performance for all models is improved using the ATM put alone instead of the underlying asset alone. The daily EPK hedge reduces the hedge portfolio standard deviation to \$11.40 per day, which is a percentage standard deviation reduction of about 60%.

The final hedge considered uses both the ATM option and the underlying as hedging instruments. Since the option and underlying price changes are not perfectly correlated, one would expect that a two asset hedge would improve on a one asset hedge. In fact, only the CRRA and average EPK methods show a relative improvement in hedge performance, and they both remain inferior to the daily EPK method.

An analysis of the hedge portfolio weights reveals that the most effective hedge, the daily EPK hedge, recommends substantially larger proportions of the ATM put and underlying be held to hedge the OTM put than the CRRA hedge. For example, the average number of units of the ATM put to be purchased per OTM put written is .18 for the CRRA hedge and .38 for the daily EPK hedge. The average EPK hedge also has higher hedge ratios than the CRRA hedge.

The larger EPK hedge ratios are due to two factors. First, the EPK model includes a stochastic volatility factor that affects the correlation of the underlying asset price and put prices. For example, when the underlying asset price falls, volatility increases which has a positive non-linear common effect on put prices. Second, the EPK places greater weights (via the EPK) on payoffs in large negative return states which further increases ATM put, OTM put, and underlying price correlations.

Thus, one element of misspecification of the CRRA pricing kernel is that it underestimates the

value of payoffs in large negative return states. Both the daily EPK and the average EPK incorporate the higher value placed on these payoffs by investors.

These results also indicate that the average empirical pricing kernel is misspecified, because it does not capture the time-variation in risk aversion and state probabilities embedded in option price dynamics. The daily EPK hedge incorporates the current state of investor risk aversion, while the average EPK hedge incorporates the average state of investor risk aversion. The average hedge ratios for both models are very similar (for the single asset hedges), but the performance of the daily EPK hedge is significantly better than the average EPK hedge. This indicates that the time varying risk aversion measured by the daily EPK — and used in hedge ratio calculation — contains useful information.

V. Conclusions

This paper investigates the empirical characteristics of investor risk aversion over equity return states by estimating a daily semi-parametric pricing kernel — the empirical pricing kernel (EPK). The EPK is estimated using state prices implied by option premia and a dynamic state probability model.

Evidence presented in this paper suggests that a constant relative risk aversion pricing kernel with a fixed risk aversion coefficient is an inadequate approximation of investor risk aversion over S&P500 return states, underestimating the value of payoffs in large negative return states. In addition, an average empirical pricing kernel fails to incorporate time-variation in relative risk aversion which is reflected in option price dynamics.

An option hedging methodology is developed that is compatible with time-varying risk aversion and is utilized in a test of the information in the empirical pricing kernel. Hedging performance results show that the empirical pricing kernel - asymmetric GARCH state probability hedge is superior to a constant relative risk aversion - lognormal state probability and an average empirical pricing kernel - average state probability hedge. This suggests that the empirical pricing kernel incorporates relevant information with respect to time-varying risk-aversion.

Appendix A

Derivation of the constant relative risk aversion pricing kernel in a Black-Scholes (1973) economy

(A.1) State price density:

$$y(S_T; S_t, r_f, \mathbf{s}) = e^{-r_f(T-t)} \frac{1}{S_T \mathbf{s} \sqrt{2\mathbf{p}(T-t)}} \exp\left\{-\frac{1}{2}[d]^2\right\} \quad d = \frac{\ln(S_T / S_t) + (r_f - 5\mathbf{s}^2)(T-t)}{\mathbf{s}\sqrt{T-t}}$$

Breeden and Litzenberger (1978)

(A.2) State probability density for a geometric brownian motion:

(e.g. Ingersoll, 1987, p. 312)

$$f(S_T; S_t, \mathbf{m}, \mathbf{s}) = \frac{1}{S_T \mathbf{s} \sqrt{2\mathbf{p}(T-t)}} \exp\left\{-\frac{1}{2}[c]^2\right\} \quad c = \frac{\ln(S_T / S_t) - (\mathbf{m} - 5\mathbf{s}^2)(T-t)}{\mathbf{s}\sqrt{T-t}}$$

(A.3) CRRA pricing kernel:

$$K(S_T; S_t, r_f, \mathbf{m}, \mathbf{s}) = y(S_T; S_t, r, \mathbf{s}) / f(S_T; S_t, \mathbf{m}, \mathbf{s})$$

$$K(S_T; S_t, r_f, \mathbf{m}, \mathbf{s}) = C(S_T / S_t)^{-\mathbf{g}} \quad \mathbf{g}_{CRRA} = (\mathbf{m} r_f) / \mathbf{s}^2 \quad C = .5(T-t)(\mathbf{m}^2 - \mathbf{m}\mathbf{s}^2 - r_f^2 - r_f \mathbf{s}^2) / \mathbf{s}^2$$

(A.4) Discretized CRRA pricing kernel

$$K(s_1) = \frac{e^{-r_f(T-t)} \text{Prob}^*(S_T \leq \underline{K})}{\text{Prob}(S_T \leq \underline{K})} \quad \underline{K} = (1 + -.10) * S_t$$

$$K(s_i) = \frac{e^{-r_f(T-t)} [\text{Prob}^*(S_T \leq \bar{K}) - \text{Prob}^*(S_T \leq \underline{K})]}{\text{Prob}(S_T \leq \bar{K}) - \text{Prob}(S_T \leq \underline{K})} \quad \left. \begin{array}{l} i=1..99 \\ \bar{K} = (1 + s_{u,i}) * S_t; \quad s_{u,i} = \text{return state upper bound} \\ \underline{K} = (1 + s_{l,i}) * S_t; \quad s_{l,i} = \text{return state lower bound} \end{array} \right\}$$

$$K(s_{100}) = \frac{e^{-r_f(T-t)} [1 - \text{Prob}^*(S_T \leq \bar{K})]}{1 - \text{Prob}(S_T \leq \bar{K})} \quad \bar{K} = (1 + .05) * S_t$$

$$\text{Prob}^*(S_T \leq K) = 1 - N(d^*) \quad \text{Prob}(S_T \leq K) = 1 - N(c^*)$$

$$d^* = \frac{\ln(S_T / K) + (r_f - 5\mathbf{s}^2)(T-t)}{\mathbf{s}\sqrt{T-t}} \quad c^* = \frac{\ln(S_T / K) + (\mathbf{m} - 5\mathbf{s}^2)(T-t)}{\mathbf{s}\sqrt{T-t}}$$

Bibliography

- Ait-Sahalia, Yacine and Andrew Lo, 1997, Nonparametric Risk Management and Implied Risk Aversion, Manuscript.
- Ait-Sahalia, Yacine and Andrew W. Lo, 1998, Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices, *Journal of Finance*, 53, 499-547.
- Baillie, R. T. and Tim Bollerslev, 1992, Prediction In Dynamic Models With Time-Dependent Conditional Variances, *Journal of Econometrics*, 52, 91-113.
- Bansal, R. and S. Viswanathan, 1993, No Arbitrage and Arbitrage Pricing - A New Approach, *Journal of Finance*, 48, 1231-1262.
- Bates, David, 1995, Hedging the Smile, Manuscript.
- Black, F., 1976, Studies of Stock Price Volatility Changes, *Proceedings of the Business and Economic Statistics Section*, , 177-181.
- Black, Fischer and Myron Scholes, 1973, The Pricing of Options and Corporate Liabilities, *Journal of Political Economy*, 81, 637-654.
- Bollerslev, Tim, 1987, A Conditionally Heteroskedastic Time Series Model for Speculative Prices and Rates of Return, *Review of Economics and Statistics*, , 542-547.
- Bollerslev, Tim, 1986, Generalized Autoregressive Conditional Heteroskedasticity, *Journal of Econometrics*, 31, 307-327.
- Bollerslev, Tim, Ray Y. Chou and Kenneth F. Kroner, 1992, ARCH Modeling in Finance - A Review of the Theory and Empirical Evidence, *Journal of Econometrics*, 52, 5-59.
- Bollerslev, Tim, Robert F. Engle and Daniel B. Nelson. ARCH Models. In *Handbook of Econometrics*, edited by R. F. Engle and D. L. McFadden, 2959-3038. Amsterdam: North-Holland, 1994.
- Bollerslev, Tim and Jeffrey M. Wooldridge, 1992, Quasi-Maximum Likelihood Estimation and Inference in Dynamic Models with Time Varying Covariances, *Economics Letters*, 12, 143-172.
- Breeden, Douglas T., Michael R. Gibbons and Robert H. Litzenberger, 1989, Empirical Tests of the Consumption-Oriented CAPM, *Journal of Finance*, 44, 231-262.
- Breeden, Douglas T. and Robert H. Litzenberger, 1978, Prices of State-Contingent Claims Implicit in Option Prices, *Journal of Business*, 51, 621-651.
- Brown, David P. and Michael R. Gibbons, 1985, A Simple Econometric Approach for Utility-Based Asset Pricing Models, *Journal of Finance*, 40, 359-381.

- Campbell, John, 1996, Consumption and the Stock Market: Interpreting the International Experience, NBER Working Paper, 5610.
- Chapman, David, 1997, Approximating the Pricing Kernel, *Journal of Finance*, 52, 1383-1410.
- de Boor, Carl, A Practical Guide to Splines, New York: Springer-Verlag, 1978.
- Diebold, Francis X. and Roberto S. Mariano, 1995, Comparing Predictive Accuracy, *Journal of Business and Economic Statistics*, 13, 253-263.
- Engle, Robert F., 1982, Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation, *Econometrica*, 50, 987-1007.
- Engle, R. F. and G. Gonzalez-Rivera, 1991, Semiparametric ARCH Models, *Journal of Business and Economic Statistics*, 9, 345-359.
- Engle, R. F. and Gary Lee, 1993, A Permanent and Transitory Component Model of Stock Return Volatility, UCSD discussion paper, 92-44R.
- Engle, Robert F. and Joshua V. Rosenberg, 1995, GARCH Gamma, *Journal of Derivatives*, 2, 47-59.
- Ermini, Luigi, 1989, Some New Evidence on the Timing of Consumption Decisions and Their Generating Process, *Review of Economics and Statistics*, 71, 643-650.
- Ferson, Wayne E. and Campbell R. Harvey, 1992, Seasonality and Consumption-Based Asset Pricing, *Journal of Finance*, 47, 511-552.
- 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.
- Ghysels, Eric, Andrew C. Harvey and Eric Renault. Stochastic Volatility. In *Handbook of Statistics*, edited by G. S. Maddala and C. R. Rao, 119-191. Amsterdam: North-Holland, 1996.
- Glosten, L. R., R. Jagannathan and D. E. Runkle, 1993, On the Relation Between the Expected Value and the Volatility of the Nominal Excess Return on Stocks, *Journal of Finance*, 48, 1779-1801.
- Hakansson, Nils H., 1976, The Purchasing Power Fund: A New Kind of Financial Intermediary, *Financial Analysts Journal*, 32, 49-59.
- Hansen, Lars Peter, 1982, Large Sample Properties of Generalized Method of Moments Estimators, *Econometrica*, 50, 1029-1054.
- Hansen, Lars Peter and Ravi Jagannathan, 1991, Implications of Security Market Data for Models of Dynamic Economies, *Journal of Political Economy*, 99, 225-262.
- Hansen, Lars Peter and Kenneth J. Singleton, 1983, Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns, *Journal of Political Economy*, 91, 249-265.

- Jackwerth, Jens Carlsen, 1997, Recovering Risk Aversion from Option Prices and Realized Returns, Manuscript.
- Kim, Sangjoon, Neil Shephard and Siddhartha Chib, 1998, Stochastic Volatility: Likelihood Inference and Comparison with ARCH models, *Review of Economic Studies*, 65, 361-393.
- Ljung, G. M. and G. E. P. Box, 1978, On a Measure of Lack of Fit in Time Series Models, *Biometrika*, 65, 297-303.
- Nelson, Daniel B., 1990, ARCH Models as Diffusion Approximations, *Journal of Econometrics*, 45, 7-38.
- Newey, Whitney and Kenneth West, 1987, A Simple Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix, *Econometrica*, 55, 703-708.
- Rosenberg, Joshua, 1998, Empirical Tests of Interest Rate Model Pricing Kernels, Manuscript.
- Rubinstein, Mark, 1994, Implied Binomial Trees, *Journal of Finance*, 49, 771-818.
- Shephard, Neil. Statistical Aspects of ARCH and Stochastic Volatility. In *Time Series Models in Econometrics, Finance, and Other Fields*, edited by D. R. Cox, O. E. Barndorff-Neilson and D. V. Hinkley, 1-67. London: Chapman and Hall, 1996.
- Slesnick, Daniel T., 1998, Are Our Data Relevant to the Theory? The Case of Aggregate Consumption, *Journal of the American Statistical Association*, 16, 52-61.
- Taylor, Stephen J., 1994, Modeling Stochastic Volatility: A Review and Comparative Study, *Mathematical Finance*, 4, 183-204.
- Wilcox, David, 1992, The Construction of U.S. Consumption Data: Some Facts and Their Implications for Empirical Work, *American Economic Review*, 82, 922-941.
- Zakoian, Jean-Michel, 1994, Threshold Heteroskedastic Models, *Journal of Economic Dynamics and Control*, 18, 931-955.

Table 1 - Summary of option data

Data used in estimation of state price densities (1991-1995)

Number of estimation dates = 1235, Total number of observations = 18,537

	Mean	Std. dev.	Min	Max
Number of options for state price estimation (per day)	15.01	2.85	9	24
Time until expiration	20.89	6.77	10	53
Implied standard deviation (annualized)	0.1565	0.0598	0.0676	0.8778
Moneyness range (per day)	0.1184	0.0635	0.0423	0.5350

Data used in option hedging tests (1991-1995)

*Only dates for which one day ATM and OTM price changes exist are included

	ATM put option	OTM put option
Number of observations	912	912
Average price	5.74	1.95
Std. dev. price	1.46	1.00
Average price change	-0.21	-0.10
Std. dev. price change	1.16	0.48
Average time to maturity (days)	20.33	20.33
Std. dev. time to maturity (days)	6.64	6.64
Average moneyness	0.000	-0.036
Std. moneyness	0.004	0.008
Average implied standard deviation (annualized)	0.1223	0.1526
Std. dev. of implied standard deviation	0.0263	0.0255

Table 1 summarizes the characteristics of the S&P500 index option data used in this study. On each day, one quote for each option series is obtained by averaging the midpoint of the last bid and ask quote of the day. The intraday Black-Scholes implied volatility is calculated using the synchronized S&P500 index price measured using the last transactions price for the corresponding S&P500 futures contract and the cost-of-carry relationship, the riskless rate measured using the three month treasury bill yield, the dividend yield measured using the implied dividend yield from the put-call parity formula, and the time until expiration measured using the number of trading days until expiration.

The first panel describes the data used in estimation of the state prices densities. There are a total of 1517 trading days over the period from 1991-1995. Of these, the 1235 days on which option series with at least 10 but no more than 60 days until expiration are used in the study. The moneyness range measures daily sum of the moneyness of the deepest out-of-the-money put and the deepest out-of-the money call $|E/S_t - 1| + |S_t/E - 1|$. Time to maturity is measured in trading days, and implied standard deviations are reported in annualized terms.

The second panel describes the data used in the hedging tests. For a trading day to be included, there must be a recorded price quote for the at-the-money (within 1% of the money) and out-of-the money (at least 3% out of the money) put options on the trading day and the subsequent day. There is option data that meets these criteria on 912 days. The synchronized end-of-day option price quotes are obtained by repricing the options using the dividend-adjusted Black-Scholes formula evaluated at the intraday implied volatility and the end-of-day S&P500 index level. Option price are reported in points (\$100 units) and moneyness is $S_t/E - 1$ for calls and $E/S_t - 1$ for puts where E is the option exercise price and S_t is the closing S&P500 index level.

Table 2 - Dynamic state probability models

S&P500 daily log-return summary statistics, 1990-1995

Number of observations	Annualized mean	Annualized std. dev.	Skewness	Excess kurtosis	Normality test p-value	Serial correlation test p-value	ARCH test p-value
1517	0.0966	0.1145	-0.10	2.41	<.0001	0.0375	<.0001

Estimated state probability models

	ARCH(1)		GARCH(1,1)		Asymmetric GARCH(1,1)	
	Coefficient	Robust t-statistic	Coefficient	Robust t-statistic	Coefficient	Robust t-statistic
μ	0.0003	2.00	0.0003	1.48	0.0004	2.35
ω	4.43E-05	15.49	6.44E-07	1.37	9.23E-07	1.34
α	0.1568	3.37	0.0855	3.98	0.0167	0.17
β			0.8857	30.23	0.9366	54.09
γ					0.0526	2.24

Log-likelihood 5346.29 5386.15 5429.63

ARCH p-value <.0001 0.4205 0.9200

LR test p-value <.0001 <.0001 <.0001

* the asymmetric GARCH components model does not improve on the asymmetric GARCH(1,1) model

Summary statistics for estimated innovation density

Number of observations	Mean	Std. dev.	Skewness	Excess kurtosis	Normality test p-value	Serial correlation test p-value	ARCH test p-value
1517	-0.0027	1.0058	-0.27	2.29	<.0001	0.2152	0.9200

The first panel of Table 2 reports summary statistics for S&P500 returns over the period from 1990-1995. The normality test p-value is the p-value of the Jarque-Bera (1980) normality test statistic which measures the closeness of the empirical S&P500 log-return density to a normal density. The serial correlation p-value is the p-value of the Ljung-Box Q-statistic (1978) which measures serial correlation in the residuals using ten lagged values. The ARCH test p-value is the p-value of the Engle (1982) ARCH LM statistic which measures the presence of stochastic volatility as represented by persistence in return magnitudes. Ten return lags are used in this test. The LR (likelihood ratio) test p-value measures statistical significance of the improvement of the GARCH(1,1) model versus the ARCH model and the GARCH(1,1) with leverage model versus the GARCH(1,1) model. The reported p-value is the p-value for twice the difference between the log-likelihood of the unrestricted and restricted model. Under the null that the added variable in the unrestricted model is insignificant, the statistic will be a chi-squared variate with 1 degree of freedom— which represents the difference in the number of variables in the restricted versus unrestricted model.

The second panel reports estimation results for three nested GARCH model specifications. Estimation is accomplished by maximization of the log-likelihood function using daily log returns for the S&P500 index over the period from 1990-1995. Robust t-statistics are calculated according to the method of Bollerslev and Wooldridge (1992). The asymmetric GARCH model is defined in equations (9)-(10). The ARCH p-value is the p-value of the Engle (1982) ARCH LM statistic which measures unexplained stochastic volatility effects in the model standardized residuals, and is interpreted as a model specification test. Ten return lags are used in this test.

The third panel reports the properties of the standardized residuals from the asymmetric GARCH model. The standardized residuals are calculated by dividing the ordinary residual ϵ_i by its conditional standard deviation σ_i , and are used to estimate the empirical innovation density. Deviations of the empirical innovation density from normality are illustrated by the negative skewness and positive excess kurtosis statistics.

Table 3 - Analysis of the daily EPK relative risk aversion parameter

Properties of the daily EPK relative risk aversion parameter (1991-1995)

Number of obs.	Mean	Std. dev.	Min	Max
1235	22.14	8.91	1.21	57.99

Autocorrelations

Lag 1	Lag 2	Lag 3	Lag 4	Lag 5
0.8020	0.6990	0.5960	0.5330	0.4660

Partial autocorrelations

Lag 1	Lag 2	Lag 3	Lag 4	Lag 5
0.8020	0.1560	-0.0070	0.0630	-0.0040

Correlation with other variables

Variable	Correlation	P-value
S&P500 return	-0.0061	0.8303
Change in implied vol. - objective vol. spread	0.0689	0.0166
Change in three month Treasury Bill yield	0.0411	0.1492

Table 3 presents an analysis of properties of the daily EPK relative risk aversion parameter ($\gamma_{EPK,t}$) which summarizes the risk aversion as measured by the empirical pricing kernel over the period 1991-1995. Estimation of $\gamma_{EPK,t}$ is accomplished by minimizing the distance between a power pricing kernel with shape parameter $\gamma_{EPK,t}$ and the empirical pricing kernel. The first panel contains the summary statistics for $\gamma_{EPK,t}$. The second and third panel report measures of time-series persistence using autocorrelations and partial autocorrelations. The fourth panel reports correlation coefficients and their p-values for $\gamma_{EPK,t}$ with several variables. The implied volatility - objective volatility spread is calculated by taking the difference of the at-the-money implied volatility (calculated as the square root of the average implied variance of nearest maturity S&P500 options within one percent of the money) and the objective volatility (calculated as the conditional volatility forecast using the estimated asymmetric GARCH(1,1) model).

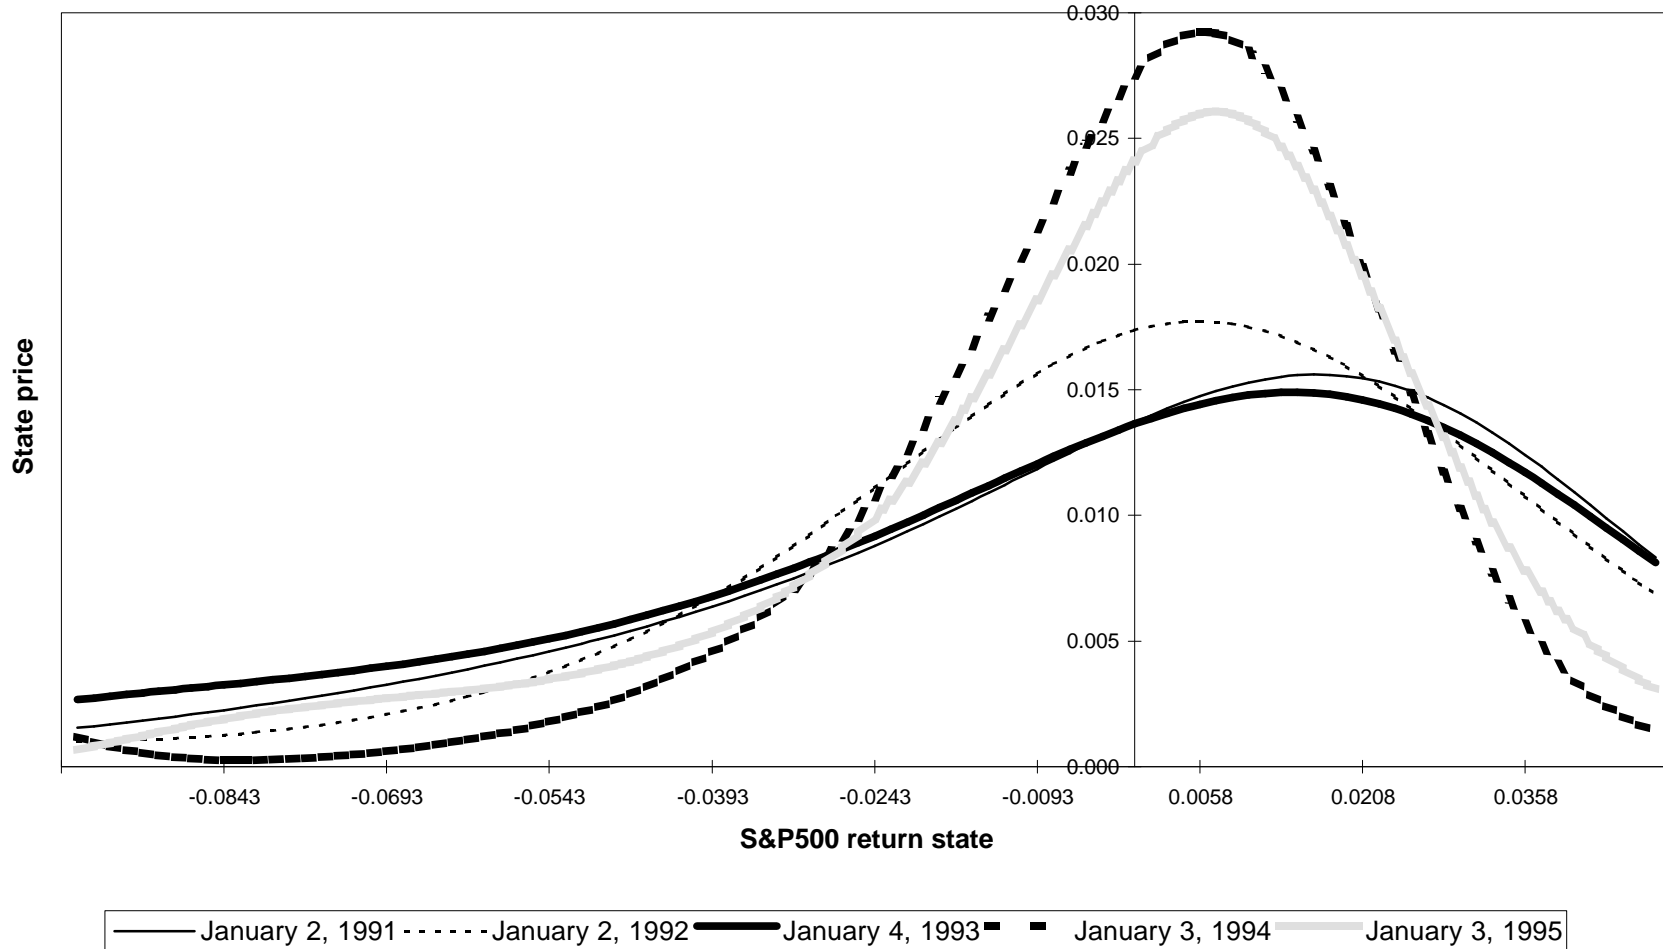
Table 4 - Hedging test results

Portfolios	Standard deviation 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	Robust t-statistic (performance of daily EPK hedge versus alternative model)
<i>No hedge:</i>				
\$100 OTM written put position	27.00	0.000	0.000	
<i>Hedge using underlying:</i>				
CRRA - lognormal hedge	20.37	-0.074	0.000	6.80
Average EPK - average state probability hedge	18.10	-0.208	0.000	3.47
Daily EPK - conditional state probability hedge	16.83	-0.191	0.000	0.00
<i>Hedge using ATM put:</i>				
CRRA - lognormal hedge	18.04	0.000	0.176	13.46
Average EPK - average state probability hedge	12.28	0.000	0.383	3.80
Daily EPK - conditional state probability hedge	11.40	0.000	0.380	0.00
<i>Hedge using underlying and ATM put:</i>				
CRRA - lognormal hedge	16.42	0.086	0.378	11.95
Average EPK - average state probability hedge	11.67	0.102	0.574	1.35
Daily EPK - conditional state probability hedge	11.54	-0.002	0.376	0.00

Table 4 reports the hedging test results for the 912 sample dates over the period 1991 - 1995. The particular hedging problem chosen is hedging a \$100 position in out-of-the-money (OTM) S&P500 index put options using at-the-money (ATM) put options, the S&P500 index portfolio, or both. Hedge ratios are constructed to neutralize the hedge portfolio sensitivity to the first-order (and in one case, second order) effects of underlying price changes. Three pricing kernel - state probability models are used to calculate the hedge ratios: CRRA - lognormal state probability, average EPK - average state probability, and daily EPK - conditional state probability. The average hedge portfolio weights are reported in the second two columns of the table.

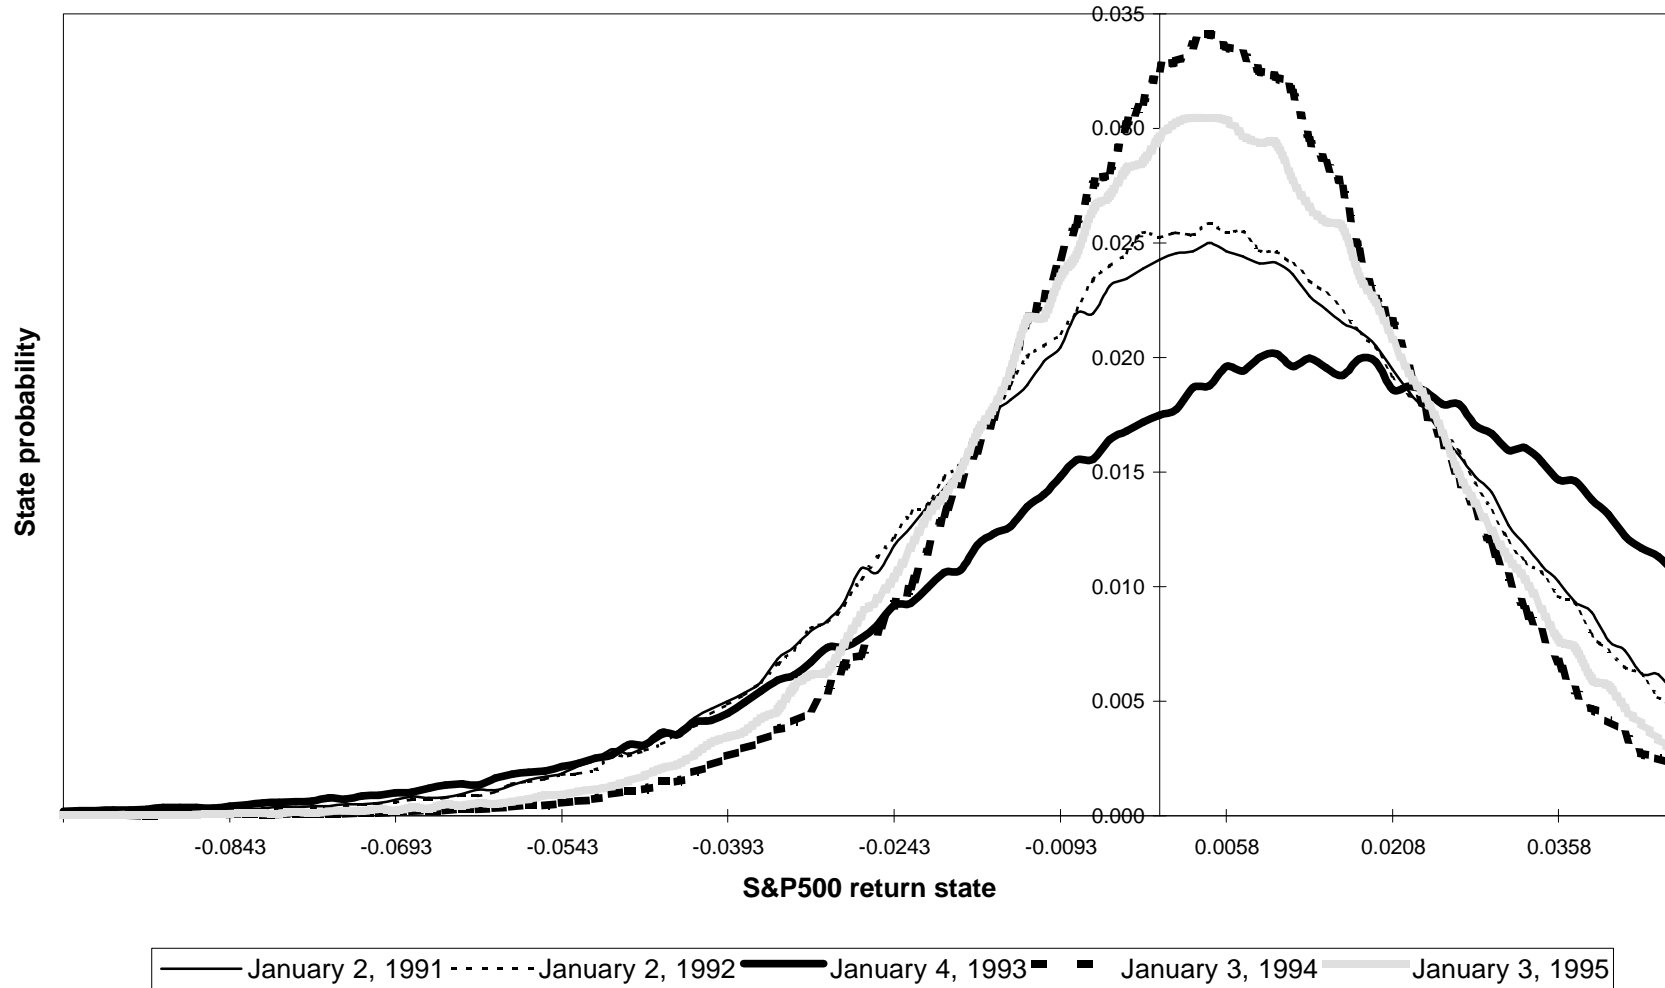
A test of the statistical significance of the improvement in hedging performance using the EPK hedge is performed using a predictive accuracy statistic similar to that proposed by Diebold and Mariano (1995). The magnitude of the hedging error of two models is compared using the t-statistic for the absolute error loss, which is defined as the difference between the absolute value of the alternative model hedging error and the absolute value of the EPK hedging error. Standard errors are calculated using the heteroskedasticity and autocorrelation consistent covariance matrix Newey and West (1987). When the t-statistic is significant and positive, the daily EPK hedge performance is significantly better than the alternative model hedge performance.

Figure 1
Estimated state price densities
First trading day of the year (1991-1995)



The discrete state price densities are estimated each day using the contemporaneous cross-section of nearest maturity S&P500 option premia. A semi-parametric state price density estimation technique is utilized which is based on the relationship between the state price density and derivatives of the call option pricing formula with respect to the option exercise price. States are defined by S&P500 index returns.

Figure 2
Estimated state probability densities
First trading day of the year (1991-1995)



The discrete state probability densities are obtained each day using monte-carlo simulation (200,000 replications) of the estimated asymmetric GARCH model with an empirical innovation density. Estimation details are given in Table 2.

Figure 3
Annualized S&P500 conditional volatility forecasts
Asymmetric GARCH model
1990 - 1995


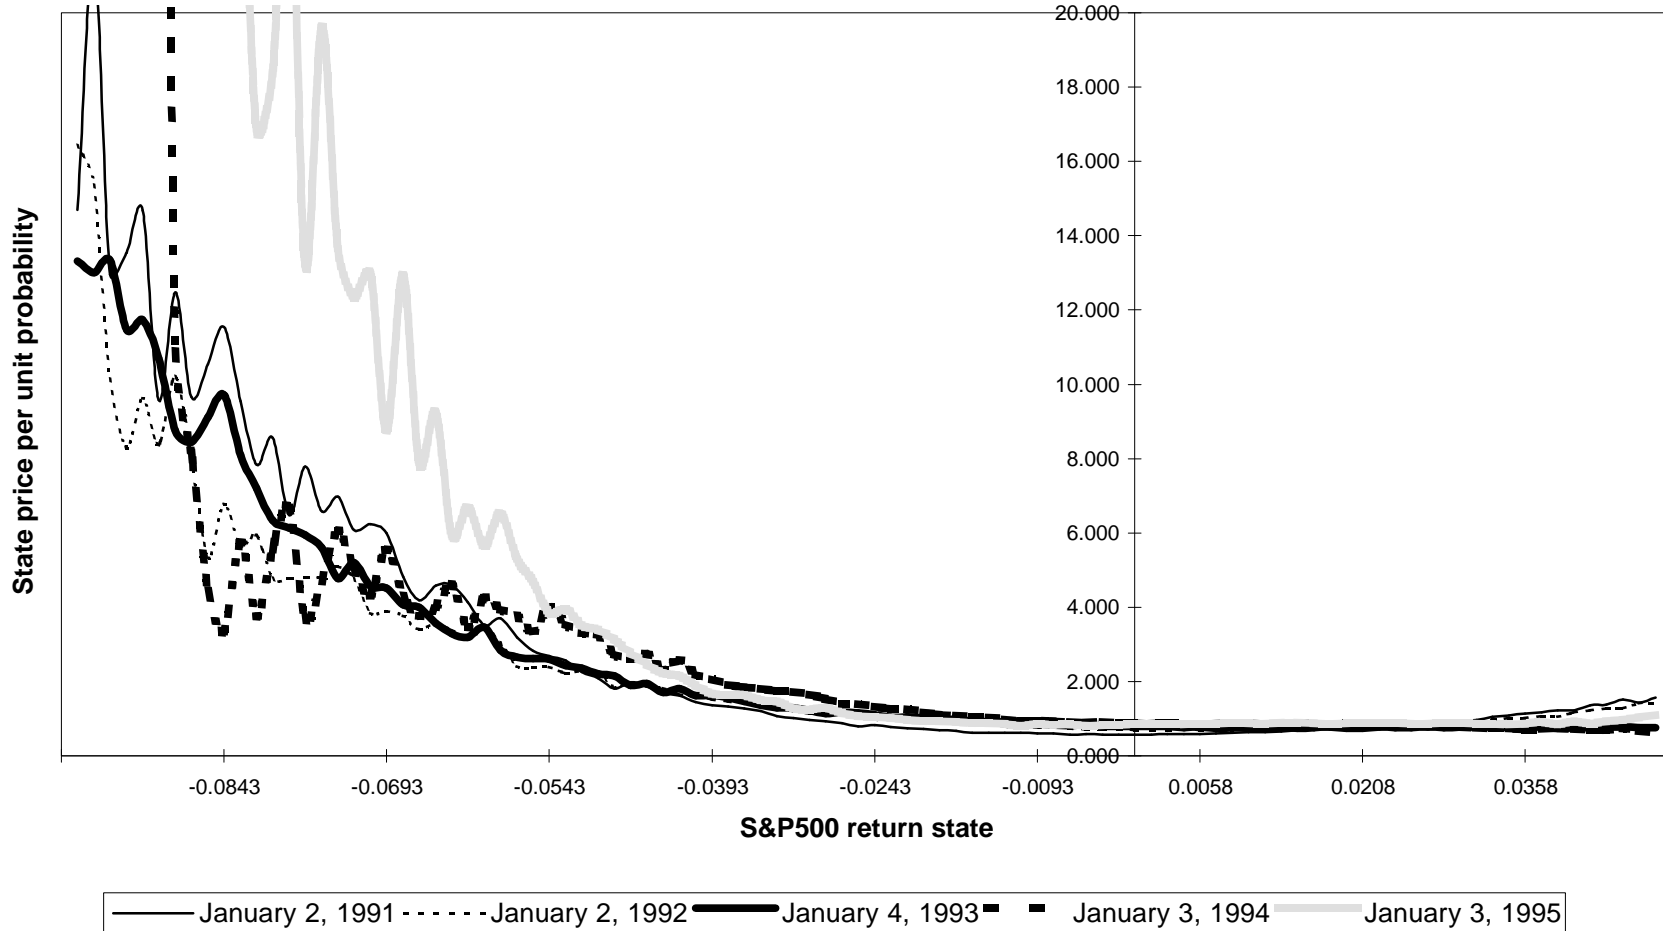


Figure 3 plots the annualized conditional standard deviation from the estimated asymmetric GARCH model defined in equations (9) - (10).

Figure 4
Empirical pricing kernels
First trading day of the year (1991-1995)



The empirical pricing kernels are estimated each day using the ratio of the daily discrete state price density and the daily discrete state probability density. Using S&P500 returns as a proxy for aggregate consumption, the negative slope of the empirical pricing kernels indicates that investors value unit payoffs in low consumption states more than in high consumption states. The apparent time variation in the empirical pricing kernel suggests that constant risk aversion over equity return states is not supported by the data.

Figure 5
Comparison pricing kernels




Figure 5 plots the estimated CRRA pricing kernel and the time-averaged empirical pricing kernel. The constant relative risk aversion pricing kernel underestimates the value of payoffs in large negative return states as indicated by its smaller negative slope.

Figure 6
Empirical pricing kernel risk aversion
1991-1995




Figure 6 presents a daily summary measure of risk aversion based on the EPK. EPK risk aversion $\gamma_{EPK,t}$ is estimated as the exponent of a power function obtained by minimizing the distance between the power function and the daily EPK.