

Empirical pricing kernels

July 2000

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 time-varying pricing kernel, which is referred to as the empirical pricing kernel (EPK). The empirical pricing kernel is the preference function that rationalizes a contemporaneous cross-section of asset prices, given a forecast payoff probability density.

We estimate the EPK on a monthly basis from 1991 to 1995 using S&P500 index option data and a stochastic volatility model for the S&P500 return process. We find substantial evidence of time-varying risk aversion over S&P500 return states. In addition, we find that empirical risk aversion over S&P500 return states is linked to business conditions; the level of risk aversion is positively correlated with indicators of recession and negatively correlated with indicators of expansion.

An option hedging methodology is developed to test the predictive information in the empirical pricing kernel and its associated state probability model. Hedging performance is significantly improved using hedge ratios based on a time-varying pricing kernel rather than a time-invariant pricing kernel.

This paper has benefited from the suggestions of Allan Timmermann, Bruce Lehmann, Steve Figlewski, Joel Hasbrouck, Ravi Bansal, Gur Huberman, Peter Carr, Jose Lopez, David Bates, Mark Broadie, 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, two anonymous referees at the Journal of Financial Economics, and William Schwert. The author wishes to thank David Hait for research assistance. The original title for this paper (first draft: April 1997) was *Option hedging using empirical pricing kernels*. The most recent version of this paper may be downloaded from:
http://www.stern.nyu.edu/~jrosenb0/J_wpaper.htm.

I. Introduction

The asset pricing kernel summarizes investor preferences for payoffs over different states of the world. In the absence of arbitrage, all asset prices can be expressed as the expected value of the product of the pricing kernel and the asset payoff. Thus, the pricing kernel, coupled with a probability model for the states, gives a complete description of asset prices, expected returns, and risk premia.¹

One pricing kernel estimation strategy is to infer a pricing kernel that rationalizes observed asset prices or returns. Estimation is considerably simplified if the pricing kernel state variables and functional form are prespecified. Hansen and Singleton (1982, 1983) postulate that aggregate U.S. consumption is the pricing kernel state variable and that the pricing kernel is a power function. Using this specification and historical consumption and return data, they estimate the two parameters that determine the pricing kernel.

Generalizations of this approach include Bansal and Viswanathan (1993), in which equity returns and an interest rate are used as pricing kernel state variables. Bansal and Viswanathan (1993) estimate the pricing kernel semi-nonparametrically using neural networks. Chapman (1997) uses functions of consumption and its lags as pricing kernel state variables, and he estimates the pricing kernel function using an orthogonal polynomial expansion. 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.

Unfortunately, consumption-based approaches for pricing kernel estimation are subject to problems with imprecise measurement of aggregate consumption. Recently, Ait-Sahalia and Lo (2000) have used option data and historical returns data to non-parametrically estimate the functional form of the pricing kernel as defined over equity return states. This technique avoids the use of aggregate consumption data or a parametric pricing kernel specification. Along similar lines, Jackwerth (2000) nonparametrically estimates the “risk-aversion function” using option data and historical returns data. A significant problem with both papers is misspecification of the state probability density. In addition, both papers use an estimation technique based on time-aggregated data that introduces bias and does not permit an analysis of time-variation in the pricing kernel.

¹ Equivalent concepts are the stochastic discount factor, the state-price-per-unit-probability, and the intertemporal marginal rate of substitution. Campbell, Lo, and MacKinlay (1997, Chapter 8) provide a thorough discussion of the role of the pricing kernel in asset pricing.

In this paper, the pricing kernel is estimated as a function of contemporaneously observed asset prices and a predicted asset payoff density. In particular, the empirical pricing kernel (EPK) is the preference function that provides the “best fit” to observed asset prices, given the payoff density. By estimating the EPK at a sequence of points in time, it is possible to observe and model the dynamic structure of the pricing kernel itself. From this analysis, we obtain improved option pricing relations, hedging parameters, and a better understanding of the pattern of risk premia.

We estimate the EPK on a monthly basis from 1991 to 1995 using S&P500 index option data and a stochastic volatility model for the S&P500 return process. We find substantial evidence of time-varying risk aversion over S&P500 return states. In addition, we find that empirical risk aversion over S&P500 return states is linked to business conditions; the level of risk aversion is positively correlated with indicators of recession and negatively correlated with indicators of expansion. This is consistent with the Fama and French (1989) result that risk-premia are highest at business cycle troughs and lowest at business cycle peaks.

For example, widening credit spreads are an indicator of economic contraction. We find that empirical risk aversion increases significantly when credit spreads widen. Flattening of the interest rate term structure is associated with periods of economic contraction (see, e.g., Estrella and Hardouvelis, 1991). We find that empirical risk aversion increases when the term structure flattens.

An option hedging methodology is developed to test the predictive information in the empirical pricing kernel. We apply this methodology to the daily hedging of S&P500 out-of-the-money put options. Hedges that allow time-variation in the pricing kernel are found to outperform hedges based on time-invariant pricing kernels.

The remainder of the paper is organized as follows. Section II describes the theory and previous research related to the pricing kernel. Section III presents the empirical pricing kernel estimation technique, empirical pricing kernel specification, and hedge ratio specification. Section IV presents the data used for estimation, and section V presents the estimation results. Section VI contains the hedging test results, and Section VII concludes the paper.

II. Theory and previous research

II.a. The pricing kernel as the intertemporal marginal rate of substitution

The first-order condition associated with solution of the optimal consumption problem delineates the relationship between payoffs, preferences (the pricing kernel), probabilities, and prices. One of the most commonly utilized preference specifications is the power utility function. The power utility function has two preference parameters; ρ defines the rate of time preference, and γ defines the level of relative risk aversion. In this case, the asset pricing equation is:

$$(1) \quad P_t = E_t \left[e^{-r} (C_{t+1} / C_t)^{-\gamma} g(C_{t+1}) \right]$$

In Equation (1), P_t is the current asset price denominated in units of consumption, C_t is the current level of consumption, C_{t+1} is the unknown level of optimal consumption in one period, and $g(C_{t+1})$ is the payoff function that defines the number of units of consumption (in the next period) generated by current ownership of the asset.

Equation (1) states that the asset price is the weighted expected payoff of the asset, where the conditional expectation is denoted by $E_t[\bullet]$. In the asset pricing literature, the weighting function —which in this case is $e^{-\rho} (C_{t+1}/C_t)^{-\gamma}$ — is often referred to as the asset pricing kernel. The weights depend on the ratio of future consumption to current consumption as well as the time and risk preference parameters.

Equation (1) may be written using a more general pricing kernel, $M(C_{t+1}, C_t)$, as:

$$(2) \quad P_t = E_t \left[M(C_{t+1}, C_t) g(C_{t+1}) \right]$$

The pricing kernel in equation (2) depends on current and future consumption, and the preference parameters are implicit. As current consumption changes, the value of a unit payoff across future consumption state may fluctuate, which may in turn generate time-variation in the pricing kernel.

For time-separable utility functions, the pricing kernel is equal to the intertemporal marginal rate of substitution (IMRS): $M(C_t, C_{t+1}) = U'(C_{t+1})/U'(C_t)$. When investors obtain utility from sources other than consumption, the pricing kernel will depend on additional state variables.² For example, in the habit persistence models of Abel (1990), Constantinides (1990), or Campbell and Cochrane (1999), the pricing

² When the utility function depends on an additional state variable (X_t) and the utility function is time-separable, the pricing kernel is $U'(C_{t+1}, X_{t+1})/U'(C_t, X_t)$. See, for example, Campbell, Lo, and MacKinley (1997, Sec. 8.4).

kernel depends on past as well as current consumption. In Eichenbaum, Hansen, and Singleton (1988) the pricing kernel depends on leisure, and in Startz (1989) the pricing kernel depends on durable goods purchases. Thus, changes in the pricing kernel may be associated with state variables other than current consumption.

Standard risk-aversion measures may be expressed as the normalized slope of the pricing kernel. For example, the Arrow-Pratt measure of absolute risk-aversion is the negative of the ratio of the derivative of the pricing kernel to the pricing kernel.³ The Arrow-Pratt measure of relative risk aversion is absolute risk aversion multiplied by current consumption.

In many cases, the state variables that affect the pricing kernel also affect the level of risk aversion. Campbell (1996) shows that a habit persistence utility function exhibits time-varying relative risk aversion, where relative risk aversion is decreasing in the amount that consumption exceeds the habit (the surplus consumption ratio). Hence, we might expect to see decreases in relative risk aversion during economic expansions when consumption is high relative to the habit and an increase in relative risk aversion during economic contractions when consumption falls closer to the habit. In contrast, the power utility function exhibits (time-invariant) constant relative risk aversion.

While equation (1) is often 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 an asset's payoffs, and payoff probabilities, and price. For example, Hansen and Singleton (1982) consider 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.

$$(3) \quad 0 = E \left[e^{-r} (C_{t+1} / C_t)^{-\alpha} (P_{t+1} / P_t) - 1 \right]$$

Then, estimation of the preference parameters is accomplished using the sample analog to equation (3), 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

³ The Arrow-Pratt measure of absolute risk aversion is given by $-U''(C_{t+1})/U'(C_{t+1})$. So, the negative of the ratio of the derivative of the pricing kernel to the pricing kernel is $-(dM_t/dC_{t+1}) * (1/M_t) = (U''(C_{t+1})/U'(C_t)) * (U'(C_t)/U'(C_{t+1})) = -U''(C_{t+1})/U'(C_{t+1})$.

consumption over a quarter (or month) as measured in the National Income and Products Accounts. The gross asset return (P_{t+1}/P_t) is taken to be the observed return to a diversified U.S. equity portfolio.

A potentially significant problem with estimation techniques that use aggregate consumption data is measurement error. 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.

To mitigate the problems with measurements of aggregate consumption, some papers use the gross return to a diversified equity portfolio ($r_{t+1} = S_{t+1}/S_t$) as a proxy for the aggregate consumption growth rate (C_{t+1}/C_t).⁴ Under power utility, consumption appears in the pricing kernel in terms of the gross consumption growth rate, so $M(C_t, C_{t+1}) = e^{-p}(C_{t+1}/C_t)^{-\gamma} \cong e^{-p}(r_{t+1})^{-\gamma} = M(r_{t+1})$. In this case, the equity index return is the pricing kernel state variable.

In addition, a return-based pricing kernel is economically important, regardless of the closeness of the approximation to the consumption-based pricing kernel, for the following reasons. First, the return-based pricing kernel may be interpreted as a projection of the consumption-based pricing kernel onto equity return states, so the characteristics of the return-based pricing kernel indirectly reveal the characteristics of the consumption-based pricing kernel. Second, the return-based pricing kernel provides insight into preferences of investors across a dimension of interest, namely, states defined by returns of an equity index portfolio. Third, the return-based pricing kernel reveals the relative importance of return state probabilities and state preferences in valuation of the index portfolio (and index derivatives). And, changes in the return-based pricing kernel indicate the relative importance changing probabilities and changing preferences in generating time-varying expected returns.

II.b. The pricing kernel as the state-price-per-unit-probability

⁴ Brown and Gibbons (1995) show that when consumption is a fixed fraction (k) of wealth (W_t), which is held in a diversified equity portfolio, $r_{t+1} = (1-k)(C_{t+1}/C_t)$. Notice that $C_t = kW_t$ and $C_{t+1} = kW_{t+1}$, and $W_{t+1} = r_{t+1}(1-k)W_t$. So, $C_{t+1}/C_t = (kW_{t+1})/(kW_t) = r_{t+1}(1-k)$. When consumption is a small fraction of wealth, the portfolio return is approximately equal to the gross consumption growth rate, i.e. $r_{t+1} \cong C_{t+1}/C_t$.

An alternative representation of the pricing kernel may be obtained using the state price density formulation of the asset pricing equation. The state price density defines the prices of a continuum of claims that pay one unit of consumption in one state of the world, and zero otherwise.

Let $\psi_t(z_{t+1})$ be the state price density on date t , where states of the world are indexed by next period's state vector (z_{t+1}) . Then, the price of an asset with a payoff function of $g(z_{t+1})$ is given by:⁵

$$(4) \quad P_t = \int \mathbf{y}_t(z_{t+1}) g(z_{t+1}) dz_{t+1}$$

Equation (4) shows that the value of an asset is the integral of the product of the asset's state-dependent payoff and the state price density. Since an asset is equivalent to a bundle of state contingent claims, the value of the asset is the same as the value of the bundle of state contingent claims.

Now, let $\pi_t(z_{t+1})$ be the state probability density on date t , defined over states of the world indexed by z_{t+1} . Then, rewriting equation (4):

$$(5) \quad P_t = \int \left[\frac{\mathbf{y}_t(z_{t+1})}{\mathbf{p}_t(z_{t+1})} \right] g(z_{t+1}) \mathbf{p}_t(z_{t+1}) dz_{t+1} = E_t \left[\left[\frac{\mathbf{y}_t(z_{t+1})}{\mathbf{p}_t(z_{t+1})} \right] g(z_{t+1}) \right] = E_t [M_t(z_{t+1}) g(z_{t+1})]$$

Notice that equation (5) is identical to equation (2), so the pricing kernel is equivalent to the state-price-per-unit-probability.

$$(6) \quad M(z_{t+1}) = \frac{\mathbf{y}_t(z_{t+1})}{\mathbf{p}_t(z_{t+1})}$$

The pricing kernel defined by equation (6) is potentially time-varying, since it depends on state prices and state probabilities that reflect investor's current (date t) preferences and probability beliefs. As noted in the

⁵ Notice that the state price formulation of the asset pricing equation may be obtained from equation (2) by defining the state vector z_{t+1} as next period's optimal consumption C_{t+1} . Then, the state price density is the product of the pricing kernel and the state probability density, i.e. $\psi_t(C_{t+1}) = M_t(C_t, C_{t+1}) \pi_t(C_{t+1})$, where $\pi_t(C_{t+1})$ is the probability density of consumption at date $t+1$. This definition satisfies the characteristics of a state price density, since each state price represents the value of a security that pays one unit of consumption in the state of the world when next period's consumption equals C_{t+1} and zero otherwise.

previous section, the pricing kernel will generally depend on the current level of consumption, and may depend on additional macroeconomic state variables.

Now, consider the asset pricing equation when the asset payoff depends on only one of the elements of the pricing kernel state vector. Let z_{t+1}^1 be the first element of the state vector and $g(z_{t+1}^1)$ be the payoff function for a particular asset. Then, equation (5) specializes to:

$$(7) \quad P_t = E_t \left[M_t(z_{t+1}) g(z_{t+1}^1) \right] = E_t \left[E_t \left[M_t(z_{t+1}) \middle| z_{t+1}^1 \right] g(z_{t+1}^1) \right] = E_t \left[M_t^1(z_{t+1}^1) g(z_{t+1}^1) \right]$$

Equation (7) provides a univariate pricing kernel, $M_t^1(z_{t+1}^1) = E_t \left[M_t(z_{t+1}) \middle| z_{t+1}^1 \right]$, that is sufficient for valuation of assets with payoffs that only depend on the first element of the pricing kernel state vector. This univariate pricing kernel is the conditional expectation of the full pricing kernel with respect to the first state variable.

Two papers — Ait-Sahalia and Lo (2000) and Jackwerth (2000) — have focused on estimation of a univariate pricing kernel (or risk aversion function) defined over equity return states, i.e. $z_{t+1}^1 = r_{t+1}$. Both papers use a pooled cross-section and time-series of equity option prices to non-parametrically estimate the state price density. And, both use a four-year window of historical returns to estimate the state probability density.⁶ Then, the pricing kernel is approximated using equation (7).

These papers average state prices and state probabilities over time, so their estimates are perhaps best interpreted as a measure of the average pricing kernel over the sample period. However, the proposed estimators in these papers are biased. The sample analogue for equation (7) would be the sample average of the ratio of conditional state prices and conditional state probabilities. Instead of using an average of ratios, Ait-Sahalia and Lo (2000) and Jackwerth (2000) choose to take a ratio of averages. Jensen's inequality implies that these estimators are not equivalent.

The estimation techniques of Ait-Sahalia and Lo (2000) and Jackwerth (2000) techniques are also problematic, because of the misspecification of investors beliefs about future return probabilities. Both papers use a four-year window for estimation of the state probabilities, so that investors “forget” about events more

⁶ Ait-Sahalia and Lo (2000) estimate average state prices using a pooled cross-section and time-series of S&P500 index option data over 1993. Average state probabilities are estimated by kernel-smoothing the return histogram of S&P500 index returns from 1990 – 1993. Jackwerth (2000) estimates average state prices using a pooled cross-section and time-series of

than four years old. And, events over the prior four years receive equal importance in predicting future probabilities. Evidence from the extensive literature on stochastic volatility — e.g. Bollerslev, Chou, and Kroner (1992) — suggests that neither assumption is correct. Future state probabilities depend more on the recent events than long-ago events, but long-ago events still remain useful for prediction of future state probabilities. These types of misspecification of the state probability density will result in errors in the estimated pricing kernel and risk-aversion function.

In addition, the applicability of average pricing kernels and risk-aversion functions for asset pricing and hedging is limited, since assets will only be correctly priced when risk aversion and state probabilities are at their average level. When either deviates from its average, inaccurate prices (and hedge ratios) will result. Ait-Sahalia and Lo (2000) 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. 22).”

III. Empirical pricing kernel estimation strategy

III.a. Estimation technique

This paper proposes an estimator for a time-varying return-based pricing kernel (empirical pricing kernel) as the preference function that rationalizes a contemporaneous cross-section of asset prices using forecast payoff probabilities. The EPK forecast payoff probabilities are obtained using a stochastic volatility model that captures salient features of the asset return process. In addition, the EPK is estimated using a contemporaneous cross-section of option prices, rather than a pooled cross-section and time-series of option prices. Hence, the EPK represents an estimate of the pricing kernel at a particular date, rather than an approximation of an average pricing kernel over a period of a year or longer.

Consider the case in which the pricing kernel is defined over the one-period return (r_{t+1}) for a particular asset (“the underlying asset”). Let $P_{i,t}$ be the price of the i^{th} asset whose payoff function, $g(r_{t+1})$, is solely a

S&P500 index option data over four periods: 860402-871018, 881019-910318, 910319-930818, 930119–951229. Average state probabilities are estimated by kernel-smoothing the return histogram of S&P500 index returns over four year periods.

function of the one-period underlying asset return. Let $f_t(r_{t+1})$ be the forecast probability density of one-period returns. Then, equation (1) specializes to:

$$(8) \quad P_{i,t} = E_t [M_t(r_{t+1})g_i(r_{t+1})] = \int M_t(r_{t+1})g_i(r_{t+1})f_t(r_{t+1})dr_{t+1}$$

Suppose the pricing kernel is represented parametrically as $M_t(r_{t+1}) = M_t(r_{t+1};\theta_t)$, where θ_t is an $N \times 1$ parameter vector. Then, identification of the pricing kernel requires the prices of at least N assets whose payoffs are determined by r_{t+1} and the forecast probability density given by $f_t(r_{t+1})$.

The pricing kernel parameter vector (θ_t) may be estimated by minimizing the distance between the observed asset prices ($P_{i,t}$) and a set of fitted prices ($\hat{P}_{i,t}(\mathbf{q}_t)$) obtained using equation (8). If the distance function is chosen to be the sum of squared pricing errors, then estimation is accomplished by solving a non-linear least squares problem. When there are L observed prices of assets whose payoff depends on the underlying asset return, then the empirical pricing kernel is $M_t(r_{t+1},\theta_t)$ evaluated at the solution of:

$$(9) \quad \underset{\mathbf{q}_t}{\text{Min}} \sum_{i=1}^L [P_{i,t} - \hat{P}_{i,t}(\mathbf{q}_t)]^2$$

An additional useful identifying condition is given by the valuation of a riskless one-period bond, $B_t = E_t[M_t(r_{t+1};\theta_t)]$. We include a scaling factor in each parametric pricing kernel specification, $(\theta_{0,t})$, which is selected to ensure that the riskless bond price is matched exactly.

When the payoff density, $f_t(r_{t+1})$, is represented by a set of J realized (or simulated) returns, rather than in closed form, the fitted asset price may be estimated using the following approximation to equation (8), where averaging replaces integration:

$$(10) \quad \hat{P}_{i,t}(\mathbf{q}_t) \cong J^{-1} \sum_{j=1}^J [M_t(r_{t+1,j};\mathbf{q}_t)g_i(r_{t+1,j})]$$

III.b. Pricing kernel specifications

We consider two pricing kernel specifications. In the first specification, the pricing kernel is a power function of the underlying asset gross return.

$$(11) \quad M_t(r_{t+1}; \theta_t) = \theta_{0,t}(r_{t+1})^{-\theta_{1,t}}$$

The first parameter ($\theta_{0,t}$) is a scaling factor. The second parameter ($\theta_{1,t}$) determines the level of relative risk aversion at date t . When the second parameter is time-invariant, this pricing kernel exhibits constant relative risk aversion. Otherwise, this pricing kernel only exhibits constant relative risk aversion over a single period.

We are also interested in a more general specification of the pricing kernel that permits time-varying risk aversion and a variety of shapes. Following Chapman (1997), we consider a low-order orthogonal polynomial expansion to estimate the pricing kernel.⁷ Orthogonal polynomial families have the useful property that all terms are mutually orthogonal under a particular weighting function. As noted by Chapman (1997), orthogonal polynomials mitigate estimation problems due to collinearity of a polynomial expansion, and improve estimation efficiency. These estimators also have a nonparametric interpretation, if the order of the polynomial expansion increases with the sample size.

Consider the pricing kernel defined as the exponential of an orthogonal polynomial expansion with coefficients given by $\theta_{0,t} \dots \theta_{N,t}$ and terms given by $T_0(r_{t+1}) \dots T_N(r_{t+1})$.

$$(12) \quad M_t(r_{t+1}; \theta_t) = [\theta_{0,t} T_0(r_{t+1})] \exp[\theta_{1,t} T_1(r_{t+1}) + \theta_{2,t} T_2(r_{t+1}) + \dots + \theta_{N,t} T_N(r_{t+1})]$$

We select the generalized Chebyshev polynomial expansion, which is discussed in Judd (1998, Ch. 6), that has terms defined by $T_n(x) = \cos(ncos^{-1}(x))$.⁸ In this case, $T_0 = 1$, $T_1 = x$, and higher order terms are periodic functions over the approximation interval. We utilize an exponential function of the polynomial expansion to ensure that the pricing kernel is strictly positive.

⁷ Chapman (1997) estimates the asset pricing kernel as a function of aggregate consumption using a five term Legendre orthogonal polynomial expansion. Bansal, Hsieh, and Viswanathan (1993) estimate an international asset pricing kernel as a function of the Eurodollar interest rate and a world equity index return. This asset pricing kernel is specified using a five term polynomial expansion.

III.c. State probability density specification

Estimation of $f(r_{t+1})$, the forecast probability density of one-period returns, is accomplished using a stochastic volatility model that has been shown to incorporate key features of asset return data.⁹ 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).

The empirical section of this paper focuses on forecasting the density of equity index returns. 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: (1) return volatility is stochastic and mean-reverting, (2) return volatility responds asymmetrically to positive and negative returns, and (3) return innovations are non-normal.

These important features of the equity returns process may be parsimoniously modeled within the ARCH framework. Stochastic, mean-reverting volatility dynamics may be represented using the ARCH, GARCH, or GARCH components models of Engle (1982), Bollerslev (1986), or Engle and Lee (1993). The asymmetric feature of equity return volatility, first noted by Black (1976), may be incorporated using the ARCH model extensions proposed by Glosten, Jagannathan, and Runkle (1993) and Zakoian (1994). ARCH models have also been generalized to allow for non-normal return innovations. Bollerslev (1987) introduces a GARCH model with t-distributed innovations, while Engle and Gonzales-Rivera (1991) derive a semi-parametric model of the innovation density.

In this paper, the equity index return process is specified using an asymmetric GARCH model with an empirical 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 empirical innovation density that captures the potential non-normalities of the true innovation density.

The asymmetric GARCH model is specified as follows:

⁸ The *generalized* Chebyshev polynomial has $x = ((2r_{t+1} - a - b)/(b-a))$, where a and b are the endpoints of the approximation interval. In contrast, the Chebyshev polynomial ($x = r_{t+1}$) is defined over the approximation interval $[-1,1]$.

⁹ For the state probability density estimation, r_{t+1} refers to the one-period net return $(S_{t+1}/S_t - 1)$.

$$(13) \quad \ln(S_t / S_{t-1}) - r = \mathbf{m} + \mathbf{e}_t, \quad \mathbf{e}_t \sim f(0, \mathbf{s}_{t|t-1}^2)$$

$$(14) \quad \mathbf{s}_t^2 = \mathbf{w}_1 + \mathbf{w}_2 I + \mathbf{a} \mathbf{e}_{t-1}^2 + \mathbf{b} \mathbf{s}_{t-1|t-2}^2 + \mathbf{g} \text{Max}[0, -\mathbf{e}_{t-1}]^2$$

Equation (13) 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 (f) that has time-varying stochastic variance ($\sigma_{t|t-1}^2$). While a constant expected return is not usually compatible with time-varying risk-aversion, the effect over a short time horizon (e.g. one month) is not likely to be important, and equation (13) may be viewed as an approximation.

Equation (14) states that the daily conditional return variance ($\sigma_{t|t-1}^2$) depends on two constants (ω_1 and ω_2) the one-day lagged squared innovation (ε_{t-1}^2), and whether the lagged return was negative ($\text{Max}[0, -\varepsilon_{t-1}]$). The second constant (ω_2) permits a shift in long-run volatility during the estimation period, where I is an indicator variable that takes the value zero prior to a predetermined date, and the value of one after this date.

The parameters of the stochastic volatility model are estimated by maximum likelihood under the assumption that the innovation 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 innovation density is non-normal.

The empirical innovation density (f) is defined by dividing the each return by its conditional standard deviation. This empirical density will reflect excess skewness, kurtosis, and extremal return behavior that is not explained by the stochastic volatility model.

Monte-Carlo simulation of the asymmetric GARCH model is used to obtain the conditional state probability densities. For example, suppose states of the world are defined by the asset's one-month return. If the GARCH model is estimated using daily data, then a simulated one-month (twenty-day) return is obtained by randomly sampling twenty one-day innovations ($\varepsilon_{t+1}, \varepsilon_{t+2}, \dots, \varepsilon_{t+20}$). Then, tomorrow's simulated price, S_{t+1} , is equal to $\exp(\mu + \varepsilon_{t+1} \sigma_{t+1|t})$, where S_t is the current asset price and $\sigma_{t+1|t}$ is tomorrow's conditional volatility. Next, the conditional volatility for date $t+2$, $\sigma_{t+2|t+1}$, is calculated using the simulated return on date $t+1$ and the estimated GARCH model. Simulated prices for the subsequent nineteen days are obtained by repeating this process. Finally, the simulated one-month return equals $S_{t+20}/S_t - 1$.

III.d. Hedge ratio specification

In this paper, we construct hedge ratios using a general technique that is consistent with an arbitrary specification of state probabilities and pricing kernel. These hedge ratios are designed to neutralize an option portfolio to the first and second-order effects of changes in the underlying price. In a continuous-time diffusion setting, a first-order hedge is sufficient to eliminate all randomness in the hedge portfolio and also provides a minimum-variance hedge. In a discrete-time setting with stochastic volatility, first and second-order hedges will reduce, but not eliminate hedge portfolio variability.

Using a Taylor series expansion, the put option price change ($Put_{t+1} - Put_t$) after one day is approximately equal to the following function of the underlying price change ($S_{t+1} - S_t$):

$$(15) \quad Put_{t+1} - Put_t \cong \frac{\partial Put_{t+1}}{\partial S_{t+1}} (S_{t+1} - S_t) + \frac{1}{2} \frac{\partial^2 Put_{t+1}}{\partial S_{t+1}^2} (S_{t+1} - S_t)^2$$

Hedge ratios depend on the sensitivity of the asset to be hedged to changes in the underlying price ($\partial Put_{t+1} / \partial S_{t+1}$, $\partial^2 Put_{t+1} / \partial S_{t+1}^2$). The hedge ratio is given by the negative of the ratio of the sensitivity of the instrument to be hedged and the sensitivity of the hedging instrument. For example, to eliminate the sensitivity of a portfolio containing a purchased put option to first-order underlying price effects, $-\partial Put_{t+1} / \partial S_{t+1}$ shares of the underlying asset should be purchased. To eliminate an option portfolio's sensitivity to first-order and second-order underlying price effects, two hedging instruments are used.

This paper generalizes the Engle and Rosenberg (1995) technique for estimation of option price sensitivities to a setting with an arbitrary pricing kernel. Following Engle and Rosenberg (1995), consider three possible one-day underlying price changes. The stock price could rise by one-standard deviation to $S_t + \epsilon$, remain constant at S_t , or fall by one standard deviation to $S_t - \epsilon$. Each underlying price change results in a different date $t+1$ put option price: $Put_{t+1|S_{t+\epsilon}}$, $Put_{t+1|S_t}$, or $Put_{t+1|S_{t-\epsilon}}$.

The three put prices may be used to calculate the following centered finite difference approximations to the first and second derivatives of the option pricing formula:

$$(16) \quad \frac{\partial Put_{t+1}}{\partial S_{t+1}} \cong \frac{Put_{t+1|S_{t+\epsilon}} - Put_{t+1|S_{t-\epsilon}}}{2\epsilon} \quad \frac{\partial^2 Put_{t+1}}{\partial S_{t+1}^2} \cong \frac{Put_{t+1|S_{t+\epsilon}} - 2Put_{t+1|S_t} + Put_{t+1|S_{t-\epsilon}}}{\epsilon^2}$$

The value of the put option at date $t+1$, for a given level of the underlying price (S_{t+1}), may be obtained by applying equation (8), which states that an asset's price is its pricing-kernel-weighted expected payoff. Suppose that the put option expires on date $t+T$, and the current underlying price is S_t . Using a return-based pricing kernel, with $r_{t,t+T} = S_{t+T}/S_t$ and $r_{t+1,t+T} = S_{t+T}/S_{t+1}$:

$$(17) \quad Put_t = E_t [M_t(r_{t,t+T}) \text{Max}[0, K - S_{t+T}]]$$

$$(18) \quad Put_{t+1|S_{t+1}} = E_{t+1|S_{t+1}} [M_{t+1|S_{t+1}}(r_{t+1,t+T}) \text{Max}[0, K - S_{t+T}]]$$

Evaluating equation (18) requires knowledge of the state probability density function at date $t+1$ and the pricing kernel at date $t+1$, conditional on the underlying price at date $t+1$. The conditional state probability density may be determined using the underlying price process specification, and the conditional pricing kernel depends on the dynamics of the pricing kernel.

For example, if the underlying price follows a geometric Brownian motion, then the underlying price at date $t+1$ (S_{t+1}) provides additional information about the expected price at date $t+T$ (S_{t+T}).¹⁰ If the underlying price follows an asymmetric GARCH process, then the underlying price at date $t+1$ provides additional information about the expected price as well as higher moments of the price distribution.

Under power utility, the pricing kernel as a function of future returns is approximately constant over time, i.e. $M_{t+1}(r_{t+1,t+T}) = M_t(r_{t,t+T})$ when $r_{t+1,t+T} = r_{t,t+T}$.¹¹ In general, the date $t+1$ pricing kernel will be a function, $h(\cdot)$, of the date t pricing kernel, the underlying price change, and other state variables.

$$(19) \quad M_{t+1}(r_{t+1,t+T}) = h(M_t(r_{t,t+T}), S_{t+1}, S_t, \Omega_{t+1}) \quad \Omega_{t+1} \text{ is a vector of state variables}$$

One useful specification has the parameters of today's pricing kernel (θ_t) depend linearly on a constant (α), lagged parameter values ($\theta_{0,t-1} \dots \theta_{N,t-1}$), and the one-period asset return (r_t).

¹⁰ If the underlying price follows a geometric Brownian motion with drift and diffusion parameters of μ and σ , the probability density of $\log(S_{t+T})$ conditional on S_{t+1} is $N(\ln(S_{t+1}) + (\mu - \sigma^2/2)(T-1), \sigma\sqrt{T-1})$.

¹¹ Consider a power pricing kernel defined over return states when date t and date $t+1$ are close together in time. $M_t(r_{t,t+T}) = e^{-\rho T}(r_{t,t+T})^{-\gamma}$ and $M_{t+1}(r_{t+1,t+T}) = e^{-\rho(T-1)}(r_{t+1,t+T})^{-\gamma}$. When $r_{t,t+T} = r_{t+1,t+T}$, $M_{t+1}(r_{t+1,t+T}) = e^{\rho} M_t(r_{t,t+T}) \cong M_t(r_{t,t+T})$, for small values of ρ .

$$(20) \quad \theta_{i,t} = \alpha + \beta_0 \theta_{0,t-1} + \beta_1 \theta_{1,t-1} \dots + \beta_N \theta_{N,t-1} + \gamma(r_t)$$

Then, tomorrow's pricing kernel, $M_{t+1}(r_{t+1,t+T}; \theta_{t+1})$ is estimated using the forecast of tomorrow's pricing kernel parameter vector: $\hat{q}_{i,t+1}$.

$$(21) \quad \hat{q}_{i,t+1} = E_t[\theta_{i,t+1}|r_{t+1}] = \alpha + \beta_0 \theta_{0,t} + \beta_1 \theta_{1,t} \dots + \beta_N \theta_{N,t} + \gamma(r_{t+1})$$

IV. Data

The options dataset used in this paper is obtained from a subset of Berkeley Options Database covering the period 1991 - 1995. One of the advantages of this dataset is that option quotes are time-stamped and recorded along with the simultaneously measured underlying price. This facilitates construction of a database of time-synchronized daily option "closing" prices.

The database of option closing prices is constructed using the following procedure. First, "end-of-day" option prices are collected each day for all contracts by averaging the last recorded bid-ask quote of the day between 2:00 and 3:00 PM Central Time. The cross-section of midquotes from the last hour of trading is not entirely synchronized, since the S&P500 index level may change over the last hour of trading. To correct for this, a Black-Scholes (1973) implied volatility is calculated each day for each option contract using the end-of-day midquote, the contemporaneous S&P500 index level, the appropriate riskless interest rate, time until expiration, and dividend yield.¹² The closing price for each contract is obtained by evaluating the Black-Scholes formula using the same inputs and the implied volatility, except that the closing S&P500 index level is used instead of the synchronized S&P500 level.¹³ Then, each call (put) price is averaged with the synthetic

¹² A dividend time-series for the S&P500 index is extracted from the CRSP database. The present value of dividends ($D_{i,T}$) over the life of each option contract (and the dividend yield) are calculated by discounting the realized future dividend stream using the riskless interest rates described below. The riskless interest rate is obtained as using quoted bid and ask discount rates for U.S. Treasury Bills with maturities of one, three, and six months from Datastream. The riskless rate for a particular option is calculated by linear interpolation of the interest rates of Treasury Bills that "straddle" the option expiration date. The option time-until-expiration is measured in trading days.

¹³ This technique does not require that the Black-Scholes model is correct, but simply uses the Black-Scholes formula as an extrapolation device to calculate an adjustment to the observed option price if the S&P500 index level changes in-between the time of the last quote and the close of trading. If the S&P500 level does not change, then the closing option price is identical to the end-of-day option price.

put (call) price obtained using put-call parity adjusted for dividends to obtain the closing call (put) price used in the estimation procedure.

The following screening criteria are used to eliminate data errors and ensure that closing option prices are representative of market conditions at the end of the trading day. Several of these conditions are based on Bakshi, Cao, and Chen (1997). (1) Only options with mid-quotes greater than \$3/8 and less than \$50 are included in the sample. (2) Options with annualized implied volatilities less than 5% or greater than 90% are excluded from the sample. (3) Options that do not satisfy the no-arbitrage lower bound ($P_t \geq \text{Max}[0, Ke^{-r(T-t)} - S_t + D_{t,T}]$ or $(C_t \geq \text{Max}[0, S_t - Ke^{-r(T-t)} - D_{t,T}])$) are excluded from the sample. (4) Cross-sections of calls (puts) for which the no-arbitrage condition that option premia are decreasing (increasing) in the exercise price are excluded from the sample. (5) Options that violate the maximum vertical spread premium condition ($C_t(K_1, T-t) - C_t(K_2, t) \leq K_2 - K_1$; $P_t(K_2, T-t) - P_t(K_1, t) \leq K_2 - K_1$) are excluded from the sample. (6) Only options with moneyness $-.10 \leq (K/S_t - 1) \leq .10$ are included in the sample. (7) On each trading date, for a particular exercise price and expiration date, there must be both a traded call and put that meets the preceding criteria to be included in the sample. (8) Only dates on which at least eight options with different exercise prices meeting the first eight criteria are included in the sample.

A dataset of options with one month (twenty trading days) until expiration is obtained from the preceding data using the following procedure. We first eliminate all options with greater than 24 or fewer than 16 trading days until expiration. Then, for each trading date, we select the option series with time until expiration closest to twenty days. We are left with a single cross-section of call and put options each month (around the 20th of the month) with a time-until-expiration of approximately one month. This sampling methodology is similar in spirit to Christensen and Prahbala (1998).

The first panel of Table 1 reports the properties of the one-month option contracts used in pricing kernel estimation. In the sample, there are 53 months (of 60 total) in which a one-month cross-section of options is obtained that satisfies the screening criteria. On 39 dates (one each month), there is an option series with exactly twenty days until expiration; a series with 21 days until expiration is used eight times, and the remaining six dates have option series with 18, 19, and 22 days until expiration. There is no satisfactory data in seven of the sixty months of the sample.

On a given estimation date, there are between 8 and 13 options available. There are roughly equal numbers of options with moneyness $(K/S_t - 1)$ from 3% to 0%, 0% to -3%, and -3% to -6%. There are somewhat fewer options with moneyness between -6% and -10%. The smallest number of options has

moneynesses in the range from 3% to 6%, and there are no options available with moneyness greater than 6%. On average, option contracts with higher moneyness have lower implied volatilities. This pattern is often referred to as a “volatility skew.”

Due to put-call parity, there is no additional information provided by inclusion of a call and a put with the same exercise price (or moneyness) in the estimation procedure. So, we subset the database to include only out-of-the-money put options (moneyness $\leq 0\%$) and out-of-the-money call options (moneyness $> 0\%$).

Figure 1 graphs five representative cross-sections of one-month option closing prices for June of 1991 through June of 1995 against dollar moneyness ($K-S_t$). The put premia are increasing in exercise price, and the call premia are decreasing in exercise price resulting in an inverted-V shape. The variation in the slope and height of each curve reflects differences in investor probability beliefs and risk-aversion across time.

The first panel of Table 2 reports summary statistics for the daily S&P500 index returns series (1970 - 1995) used for estimation of the state probability model. Over this period, the average annualized S&P500 index return (capital appreciation only) is 7.55%, and the annualized S&P500 return standard deviation is 14.79%. S&P500 returns exhibit negative skewness and positive kurtosis, which is consistent with a stochastic volatility model with asymmetric effects. There is also evidence of return serial correlation.

V. Estimation of the empirical pricing kernel over S&P500 return states

V.a. Estimation of S&P500 return state probability densities

The second panel of Table 2 describes the dynamic state probability model used to estimate the state probability densities. 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 2.41.

Two specification tests are used to analyze the adequacy 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 .7004. The second test — Ljung-Box (1978) Q-statistic — measures unexplained serial correlation in the standardized residuals from the estimated model. The asymmetric GARCH model passes this test with a p-value of .7378. 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 innovation density (f). The third panel of Table 2 reports characteristics of the empirical innovation density indicating that it exhibits negative skewness and positive excess kurtosis. These characteristics do not affect the conditional volatility estimates, but they do affect the higher moments of forecast state probability distribution.

The fourth panel of Table 2 compares the probability of extreme positive and negative returns using the empirical innovation distribution and a standard normal distribution. Under the standard normal assumption, innovations of magnitude greater than 5 or 10 standard deviations almost never occur (less than one time in one million). In practice, the probabilities of these extreme events is non-negligible. For example, empirical return innovations less than -5 standard deviations are observed 6 times in 10,000, while empirical return innovations less than -10 standard deviations are observed 3 times in 10,000.

Figure 2 plots the estimated one-month state probability densities for June of each year, using 200,000 Monte-Carlo simulation replications. The time-variation in state probabilities is apparent.¹⁴ For example, there are higher probabilities for large negative return states in June 1991 and June 1992 than in June of subsequent years.

Figure 3 plots the conditional volatility forecasts using the asymmetric GARCH model. Over this period, estimated annualized S&P500 volatility ranges from 6.75% to 112.86% with a standard deviation of 6.36%. The highest volatility forecasts over this period are around the time of the October 1987 market crash.

¹⁴ The annualized return standard deviation estimates (in chronological order for June 1991, June 1992, ..., June 1995) are: 15.55%, 13.03%, 11.74%, 9.98%, and 10.88%. The return density skewness estimates are -0.36, -0.43, -0.46, -0.46, and -0.47. The return density kurtosis estimates are 5.06, 5.35, 5.66, 5.54, and 5.73.

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

The empirical pricing kernel defined over one-month S&P500 return states may now be estimated using the cross-sections of one-month S&P500 option premia and the contemporaneous forecasts of the one-month S&P500 return density.¹⁵ We estimate the EPK, using equation (9), in each of the 53 months for which a cross-section of S&P500 options is available. We estimate the power pricing kernel specification and a four term orthogonal polynomial pricing kernel specification (equation 10, with $N=3$).¹⁶

Table 3 reports the estimation results. The orthogonal polynomial specification results in a much closer fit to the observed option prices than the power specification. The average forecast error standard deviation for the orthogonal polynomial specification is \$.09 with a minimum of \$.03 and a maximum of \$.24. The average forecast error standard deviation for the power specification is \$.63 with a minimum of \$.28 and a maximum of \$1.34.

Figure 4 graphs the empirical pricing kernel in June of each year (1991-1995) estimated using the power specification. The level of risk aversion varies across time, as indicated by the shifting slope of the estimated pricing kernels. The negative slope of the curves indicates that the probability normalized value of a dollar payoff is greater when the S&P500 return is low than when the S&P500 return is high. In other words, investors are risk-averse (obtain declining marginal utility) with respect to payoffs over S&P500 return states.

Figure 5 graphs the orthogonal polynomial estimates of the empirical pricing kernel in June of each year. This specification is more flexible than the power specification, and the data appears to assign greater value to large negative S&P500 return states and lesser value to large positive S&P500 return states than the power specification. The state-price-per-unit probability for large negative return states is especially volatile, potentially reflecting time-varying demand for insurance against a significant market decline. There is also some evidence that there is a region of increasing marginal utility for small positive S&P500 return states.

Average pricing kernels over the sample period are estimated by evaluating each pricing kernel specification at its average parameter estimates. The average estimates are reported in Table 4, and the average pricing kernels are graphed in Figure 6. The average orthogonal polynomial EPK has some similarities to that of Ait-Sahalia and Lo (2000, p. 36, Figure 3) in that the pricing kernel is steeply upward

¹⁵ The estimated S&P500 return density is estimated using Monte-Carlo simulation with 200,000 replications. Fitted prices are evaluated using equation (10).

sloping for large negative returns and downward sloping for large positive returns. Ait-Sahalia and Lo's pricing kernel estimate also has a region of increasing marginal utility.

Jackwerth's (2000) absolute risk aversion function is closely related to the pricing kernel and may be expressed as the negative of the ratio of the first derivative of the pricing kernel and the pricing kernel ($-M_t'(r_{t+1})/M_t(r_{t+1})$). Jackwerth (2000) notes two key empirical findings with respect to the absolute risk aversion functions: "post-crash risk-aversion functions are negative around the center [return states close to zero]... [and] risk aversion functions rise for wealth levels greater than about 0.99 [return states greater than -1%]. (p. 441)"

In the case of the power specification, neither of these statements is supported by our estimation results. Notice that for the power specification, the absolute risk aversion function is given by the exponent of the power function multiplied by the inverse of the gross return [$\theta_{1,t}*(r_{t+1})^{-1}$]. Thus, the power pricing kernel exhibits declining (but positive) absolute risk aversion as long as the exponent ($\theta_{1,t}$) is positive, which is what we find over the period 1991-1995.

In the case of the orthogonal polynomial specification, we find some evidence consistent with Jackwerth's findings. We estimate the average absolute risk-aversion function corresponding to the average orthogonal polynomial pricing kernel graphed in Figure 6. We find that there is a region of negative absolute risk-aversion over the range from -4% to 2% and that absolute risk aversion is increasing for returns greater than -4%. The shape of our estimated average absolute risk aversion function is similar to Jackwerth's (p. 442, Figure 3, Panel D).

V.c. Linkages of EPK risk aversion to business cycle conditions

The relative risk aversion function implied by the empirical pricing kernel is given by $-r_{t+1}*(M_t'(r_{t+1})/M_t(r_{t+1}))$.¹⁷ In the case of the power pricing kernel specification, the relative risk aversion function is equal to the scalar $\theta_{1,t}$. Hence, $\theta_{1,t}$ provides a summary measure of risk-aversion over one-month S&P500 return states on each estimation date. One way to analyze the dynamics of empirical risk aversion is to examine the time-series characteristics of the estimated exponent ($\theta_{1,t}$) for the power pricing kernel specification.

¹⁶ We choose an estimation interval of -10% to 10% for the EPK estimates. Outside of this range, the EPK level is set to the estimated value of the EPK at -10% or 10%.

The first panel of Table 4 provides summary statistics describing the behavior of one-month empirical risk aversion. Over the sample period, the average empirical risk aversion is 7.36. However, the level of risk aversion fluctuates substantially with a range 2.26 to 12.55. Empirical risk aversion is positively autocorrelated and is mean-reverting with a first-order auto-correlation coefficient of .45. Figure 7 plots the time-series of empirical risk aversion estimates.

It has been shown that risk premia are correlated with the business cycle (see, e.g., Fama and French, 1989); risk premia are lowest at business cycle peaks and highest at business cycle troughs. Since risk premia are determined by the level of risk (e.g. variance, skewness, and kurtosis in asset payoffs) and by the level of risk aversion, two complementary explanations of time-varying risk premia are time-varying risk and time-varying risk aversion. This paper provides evidence of time-varying risk aversion through the business cycle, which is consistent with the Fama and French (1989) results.

To identify the association of EPK risk aversion with the business cycle, several variables that are correlated with current and expected business conditions are constructed. One-month changes in business conditions indicators are calculated using changes in variables measured contemporaneously with the estimation date for each EPK and its one-month lag.

Credit (default) spreads — the difference between the yield on risky and riskless bonds — are often used as a negative indicator of economic expansion, under the rationale that default is more likely in recession than expansion. We construct the change in credit spread using the monthly percentage change in the difference between the average yield-to-maturity for Moody's Baa rated seasoned corporate bonds and the thirty-year constant maturity Treasury yield. Both data items are obtained from the Federal Reserve's H.15 release.

The slope of the yield curve has been shown by Estrella and Hardouvelis (1991) to be strongly correlated with the business cycle. Steepening of the slope is associated with expansion, and flattening of the slope is associated with contraction. We construct the change in the yield curve slope using the monthly percentage change in the difference between the thirty-year constant maturity Treasury yield and three-month constant maturity Treasury yield. These data items are also obtained from the Federal Reserve's H.15 release.

¹⁷ Ait-Sahalia and Lo (2000) use a similar definition of the risk-aversion function (p. 27, eq. 5.3). Notice that the relative risk aversion function is equal to the product of the gross return and the absolute risk aversion function.

The level of short-term interest rates is also associated with the state of the business cycle. Higher rates are associated with recession, while lower rates are associated with expansion. We use the percentage change in the three-month constant maturity Treasury yield to proxy for this indicator.

The S&P500 index level is a component of the index of leading indicators as positive indicator of expansion. We utilize the S&P500 return as business conditions indicator, measuring it as the percentage change in the S&P500 index level. In addition, the S&P500 return may also be used as a proxy for aggregate consumption growth, since they are positively correlated. A finding of negative correlation between EPK risk aversion and S&P500 returns would be consistent with decreasing relative risk aversion, since it would indicate that risk-aversion increases when consumption declines.

Two other variables are also included in the analysis. First, it is interesting to measure the persistence in EPK risk aversion by regressing EPK risk aversion on its first lag. This indicates how long risk-aversion may be expected to deviate from its long-run average. Second, it is useful to measure the association of EPK risk aversion with changes in the difference between at-the-money implied and objective volatilities. This spread is a proxy for risk-aversion based on the amount the market is pricing at-the-money options above a risk-neutral price based on the objective volatility.

We also include the aggregate U.S. consumption growth rate in the regression. This variable has been used in many studies of the stochastic discount factor. Consumption growth is calculated as the monthly percentage growth rate of per-capita non-durable goods and services (monthly, real, seasonally adjusted) from the Federal Reserve's FRED database. Monthly U.S. resident population estimates are from the Census Bureau.

All of the above independent variables are included in a multiple regression, which is reported in the second panel of Table 4. Individual univariate correlations and their p-values are also reported in the last two columns. Independent variables are measured in percent (except lagged EPK risk aversion) so that a one basis point (.01) change in the independent variable is associated with an increase in risk-aversion equal to one-hundredth of the regression coefficient.

The results reported in Table 4 indicate that empirical risk aversion varies with business conditions. One of the four indicators is significant with the correct sign in the multiple regression, two of the four indicators exhibit a significant univariate correlation with the correct sign, and one indicator has a marginally significant univariate correlation with the correct sign.

For example, credit spreads are significantly associated with risk aversion; a one-basis point widening of the credit spread is associated with an increase of .09 in the level of relative risk aversion. On a univariate basis, steeping of the term structure slope is significantly associated with lower levels of risk aversion. Increases in “riskless” interest rates are also weakly associated with higher levels of risk aversion.

Empirical risk aversion has a first-order autocorrelation coefficient of .48, indicating that deviations from the long-run average are expected to disappear after about six months. The implied - objective volatility spread is also highly significant with the correct sign. Finally, the regression results indicate that aggregate consumption growth and S&P500 returns are not correlated with empirical risk-aversion.

One of the leading alternatives to the power utility function in consumption-based asset pricing models is the habit persistence model. If S&P500 returns are considered to proxy for consumption growth rates, then our empirical results provide some support for habit persistence in the following ways.

First, we find that relative risk aversion is time-varying as implied by many habit models, rather than constant as would be the case for a power utility model. Second, we find that relative risk-aversion is counter-cyclical. This is consistent with the predictions of habit persistence models in which risk-aversion is potentially counter-cyclical (e.g. Campbell, 1996). For example, in some habit persistence models, the surplus consumption ratio (the proportion that current consumption exceeds the habit) is negatively correlated with relative risk aversion. If the surplus consumption ratio is cyclical, then these habit models would also predict that relative risk-aversion would be counter-cyclical.

VI. 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 the pricing kernel and state probability estimates, so the accuracy of each pricing kernel may be judged by measuring hedging performance using different pricing kernels.

The hedging tests are based on hedging a \$100 portfolio of 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 hedging sample is constructed using the same screening criteria as for the sample constructed for estimation of the EPK. However, the hedging sample is composed of options with approximately one-month (from 16 to 24 days) until expiration, instead of exclusively one-month options. On each sample date, we select a put with moneyness $(K/S_t - 1)$ closest to zero but no more than 1% in absolute value. This is the “at-the-money” put. We then select the put option with moneyness closest to -3% , but no greater than -3% . This is the “out-of-the-money” put. When suitable options with closing prices on the sample date and the next trading date cannot be found, the sample date is excluded from the analysis. Using these criteria, there are 243 observations available for the hedging tests. The second panel of Table 1 summarizes the sample characteristics.

We are interested in testing four different pricing kernel specifications. To this end, we construct hedge ratios based on the power and orthogonal polynomial pricing kernels. These two pricing kernel estimates are time-varying but differ in their functional form. We also construct two time-invariant pricing kernels to analyze the importance of time-variation in the pricing kernel. The time-invariant pricing kernels are simply the power and orthogonal polynomial specifications, evaluated at their average parameter estimates.

To estimate hedge ratios, it is necessary to estimate the one-day ahead pricing kernel. We initially consider the pricing kernel forecasting model given by equation (21). In this model, the one-day ahead parameter vector (θ_{t+1}) depends on the current parameter vector (θ_t) and the one-day asset return. The forecasting model is estimated using a linear regression of pricing kernel parameter estimates on one-day lagged parameter estimates and the one-day S&P500 return.

For estimation, we construct a sample of estimated pricing kernels (power and orthogonal polynomial) on each of the 243 hedging sample dates. The hedging dates that are not one day apart are dropped from the sample. The remaining 133 observations are used in the forecasting model estimation regressions. The S&P500 return is not found to be statistically significant in any of the regressions (we cannot reject $\gamma = 0$), so it is dropped from the forecasting model. Also, the scaling factor $(\theta_{0,t})$ is only included in its own forecasting regression.

Table 5 reports the estimation results. All of the parameters are reasonably well predicted by their lagged values and the lagged values of other parameters. The adjusted R-squared for the six regressions ranges from 12% to 52%. For example, the second panel of Table 5 reports the $\theta_{1,t+1}$ forecasting model regression estimates for the power pricing kernel specification. In this model, $\theta_{1,t+1}$ is forecasted by the sum of 2.26 and $.69 * \theta_{1,t}$. Both parameters (α, β_1) are statistically significant and the regression adjusted R-squared is 50%.

This model results in an average one-day ahead forecast of $\theta_{1,t+1}$ equal to 7.49 with a standard deviation of 1.87.

Hedge ratios are then obtained using the methodology described in section III.d. In particular, first and second-order exposures to underlying price changes for each pricing kernel specification are generated using the centered finite difference approximations. One-day-ahead derivative prices are obtained using the estimated pricing kernel and 200,000 simulation replications of the estimated asymmetric GARCH model.

To measure the statistical significance of the improvement in hedging performance using a hedge based on a time-varying pricing kernel (EPK) versus a time-invariant pricing kernel, we formulate a predictive accuracy statistic similar to that proposed by Diebold and Mariano (1995). The relative performance of the EPK hedge versus a time-invariant hedge is measured using the t-statistic for the squared error loss, which is defined as the average difference between the squared time-invariant kernel one-day hedging error and the squared EPK one-day hedging error. Standard errors are calculated using the heteroskedasticity and autocorrelation consistent covariance matrix of Newey and West (1987). When the t-statistic is large and positive, the EPK hedge performance is significantly better than the time-invariant pricing kernel hedge performance.

Table 6 provides the hedging test results. We find strong evidence that the pricing kernel is time-varying. In all pair-wise comparisons, the EPK hedge results in a lower hedge portfolio standard deviation than the time-invariant pricing kernel hedge. The hedging improvement provided by using a time-varying pricing kernel, which ranges from 1% to 3%, is statistically significant in three of six cases and marginally significant in one case.

For example, the second column of Table 6 reports that the unhedged \$100 OTM put position has a daily price change standard deviation of \$22.56. The most effective hedge using the ATM put alone is the EPK power specification hedge, which reduces the portfolio standard deviation to \$11.10 per day. If the time-invariant (average) power pricing kernel hedge is used, the hedge portfolio standard deviation is \$11.21. The t-statistic for the difference in performance of these two hedges is 2.82.

Hedging using the underlying asset alone somewhat diminishes hedging performance for all the models, compared to hedging using the ATM put alone. The most effective hedge using the underlying asset alone is the EPK power specification hedge, which generates a hedging error of \$12.11. The time-invariant power pricing kernel hedge generates a hedging error of \$12.41.

We also see that hedging with both the underlying asset and the ATM option results in poorer hedging performance than hedging using the ATM put alone. While it is always superior to add additional hedging instruments when the true hedge ratios are known, model error and estimation error may result in poorer hedging performance as additional hedging instruments are added.

In all hedging comparisons, the orthogonal polynomial specification results in greater hedging error than the power specification. These differences are statistically significant (not reported in Table 6) for the hedge using the underlying ($t=4.2$) and the hedge using the ATM put ($t=3.9$). In contrast to the in-sample estimation results, the out-of-sample hedging results indicate that the power specification is superior to the orthogonal polynomial specification.

VII. Conclusions

This paper proposes and implements a new estimation technique to measure time-varying investor risk aversion over equity return states. In the absence of arbitrage, the fair asset price is the pricing kernel weighted expected asset payoff. Hence, a cross-section of contemporaneous asset prices and an estimate of the payoff density function may be used to infer the asset pricing kernel. By obtaining a sequence of forecast asset payoff densities and cross-sections of asset prices, we are able to construct a time-series of empirical pricing kernels and empirical risk aversion estimates.

We estimate the empirical pricing kernel and empirical risk aversion each month over the period 1991-1995 using S&P500 index options and forecasts of the S&P500 return density. In-sample estimation results and out-of-sample testing results provide strong evidence that the pricing kernel and relative risk aversion are time-varying. Our analysis of the relationship between empirical risk aversion and business cycle indicators indicates that empirical risk aversion is counter-cyclical. The level of risk aversion is positively correlated with indicators of recession and negatively correlated with indicators of expansion. This is consistent with the results of Fama and French (1989) as well as the predictions of some habit persistence models.

We analyze two parametric specifications for the asset pricing kernel defined over S&P500 return states. An orthogonal polynomial pricing kernel specification is found to provide a much closer fit to the option price data than the power specification. However, the power pricing kernel specification generates better hedging performance than the orthogonal polynomial specification. Hence, the evidence is mixed with respect to the correct functional form for the pricing kernel.

Bibliography

- Abel, A., 1990. Asset Prices Under Habit Formation and Catching Up with the Joneses. *American Economic Review* 80, 38-42.
- Ait-Sahalia, Y., Lo, A.W., 2000. Nonparametric Risk Management and Implied Risk Aversion. *Journal of Econometrics* 94, 9-51.
- Bakshi, G., Cao, C., Chen, Z., 1997. Empirical Performance of Alternative Option Pricing Models. *Journal of Finance* 52, 2003-2049.
- Bansal, R., Hsieh, D.A., Viswanathan, S., 1993. A New Approach to International Arbitrage Pricing. *Journal of Finance* 48, 1719-1747.
- Bansal, R., Viswanathan, S., 1993. No Arbitrage and Arbitrage Pricing - A New Approach. *Journal of Finance* 48, 1231-1262.
- Black, F., 1976. Studies of Stock Price Volatility Changes. *Proceedings of the 1976 Meetings of the Business and Economic Statistics Section American Statistical Association*, 177-181.
- Bollerslev, T., 1987. A Conditionally Heteroskedastic Time Series Model for Speculative Prices and Rates of Return. *Review of Economics and Statistics* 69, 542-547.
- Bollerslev, T., 1986. Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics* 31, 307-327.
- Bollerslev, T., Chou, R.Y., Kroner, K.F., 1992. ARCH Modeling in Finance - A Review of the Theory and Empirical Evidence. *Journal of Econometrics* 52, 5-59.
- Bollerslev, T., Engle, R.F., Nelson, D.B., 1994. ARCH Models. In *Handbook of Econometrics. ARCH Models*. edited by R. F. Engle and D. L. McFadden. 2959-3038. North-Holland, Amsterdam.
- Bollerslev, T., Wooldridge, J.M., 1992. Quasi-Maximum Likelihood Estimation and Inference in Dynamic Models with Time Varying Covariances. *Economics Letters* 12, 143-172.
- Breeden, D.T., Gibbons, M.R., Litzenberger, R.H., 1989. Empirical Tests of the Consumption-Oriented CAPM. *Journal of Finance* 44, 231-262.
- Brown, D.P., Gibbons, M.R., 1985. A Simple Econometric Approach for Utility-Based Asset Pricing Models. *Journal of Finance* 40, 359-381.
- Campbell, J., 1996. Consumption and the Stock Market: Interpreting International Experience. *Swedish Economic Policy Review* 3, 251-299.

- Campbell, J., Cochrane, J., 1999. By Force of Habit: A Consumption Based Explanation of Aggregate Stock Market Behavior. *Journal of Political Economy* 107, 205-251.
- Campbell, J., Lo, A., MacKinlay, C., 1997. *The Econometrics of Financial Markets*. Princeton University Press, Princeton.
- Chapman, D., 1997. Approximating the Asset Pricing Kernel. *Journal of Finance* 52, 1383-1410.
- Christensen, B.J., Prabhala, N.R., 1998. The Relation between Implied and Realized Volatility. *Journal of Financial Economics* 50, 125-150.
- Constantinides, G.M., 1990. Habit Formation: A Resolution of the Equity Premium Puzzle. *Journal of Political Economy* 98, 519-543.
- Diebold, F.X., Mariano, R.S., 1995. Comparing Predictive Accuracy. *Journal of Business and Economic Statistics* 13, 253-263.
- Eichenbaum, M.S., Hansen, L.P., Singleton, K.J., 1988. A Time Series Analysis of Representative Agent Models of Consumption and Leisure Choice Under Uncertainty. *Quarterly Journal of Economics* 103, 51-78.
- Engle, R.F., 1982. Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica* 50, 987-1007.
- Engle, R.F., Gonzalez-Rivera, G., 1991. Semiparametric ARCH Models. *Journal of Business and Economic Statistics* 9, 345-359.
- Engle, R.F., Lee, G., 1993. A Permanent and Transitory Component Model of Stock Return Volatility. UCSD discussion paper, 92-44R.
- Engle, R.F., Rosenberg, J.V., 1995. GARCH Gamma. *Journal of Derivatives* 2, 47-59.
- Ermini, L., 1989. Some New Evidence on the Timing of Consumption Decisions and Their Generating Process. *Review of Economics and Statistics* 71, 643-650.
- Estrella, A., Hardouvelis, G.A., 1991. The Term Structure as a Predictor of Real Economic Activity. *Journal of Finance* 46, 555-576.
- Fama, E.F., French, K.R., 1989. Business Conditions and Expected Returns on Stocks and Bonds. *Journal of Financial Economics* 25, 23-49.
- Ferson, W.E., Harvey, C.R., 1992. Seasonality and Consumption-Based Asset Pricing. *Journal of Finance* 47, 511-552.

- Ghysels, E., Harvey, A.C., Renault, E., 1996. Stochastic Volatility. In Handbook of Statistics. Stochastic Volatility. edited by G. S. Maddala and C. R. Rao. 119-191. North-Holland, Amsterdam.
- Glosten, L.R., Jagannathan, R., Runkle, D.E., 1993. On the Relation Between the Expected Value and the Volatility of the Nominal Excess Return on Stocks. *Journal of Finance* 48, 1779-1801.
- Hansen, L.P., Jagannathan, R., 1991. Implications of Security Market Data for Models of Dynamic Economies. *Journal of Political Economy* 99, 225-262.
- Hansen, L.P., Singleton, K.J., 1982. Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models. *Econometrica* 50, 1269-1286.
- Hansen, L.P., Singleton, K.J., 1983. Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns. *Journal of Political Economy* 91, 249-265.
- Jackwerth, J., 2000. Recovering Risk Aversion from Option Prices and Realized Returns. *Review of Financial Studies* 13, 433-451.
- Judd, K., 1998. *Numerical Methods in Economics*. MIT Press, Cambridge.
- Ljung, G.M., Box, G.E.P., 1978. On a Measure of Lack of Fit in Time Series Models. *Biometrika* 65, 297-303.
- Newey, W., West, K., 1987. A Simple Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica* 55, 703-708.
- Shephard, N., 1996. Statistical Aspects of ARCH and Stochastic Volatility. In *Time Series Models in Econometrics, Finance, and Other Fields. Statistical Aspects of ARCH and Stochastic Volatility*. edited by D. R. Cox, O. E. Barndorff-Neilson and D. V. Hinkley. 1-67. Chapman and Hall, London.
- Slesnick, D.T., 1998. Are Our Data Relevant to the Theory? The Case of Aggregate Consumption. *Journal of the American Statistical Association* 16, 52-61.
- Startz, R., 1989. The Stochastic Behavior of Durable and Non-durable Consumption. *Review of Economics and Statistics* 71, 356-363.
- Wilcox, D., 1992. The Construction of U.S. Consumption Data: Some Facts and Their Implications for Empirical Work. *American Economic Review* 82, 922-941.
- Zakoian, J.M., 1994. Threshold Heteroskedastic Models. *Journal of Economic Dynamics and Control* 18, 931-955.

Table 1 - Summary of one-month option data

Data used in estimation one-month empirical pricing kernels (1991-1995)

Number of estimation dates = 53, Minimum obs. per day = 8, Maximum obs. per day = 13

Option moneyness ($K/S_t - 1$)	Number of observations	Average implied volatility	Average call price	Average put price
-6% to -10%	66	20.85%	\$34.83	\$1.19
-3% to -6%	133	16.58%	\$21.81	\$1.92
0% to -3%	136	13.67%	\$10.76	\$4.13
0% to 3%	137	11.36%	\$3.25	\$9.77
3% to 6%	37	11.61%	\$0.90	\$16.66

Data used in one-month option hedging tests (1991-1995)

	ATM put option	OTM put option
Number of observations	243	243
Average price	\$6.34	\$2.39
Std. dev. price	\$1.04	\$0.77
Average price change	-\$0.11	-\$0.10
Std. dev. price change	\$1.38	\$0.56
Average time to maturity (days)	19.88	19.88
Std. dev. time to maturity (days)	2.40	2.40
Average moneyness ($K/S_t - 1$)	0.00%	-3.61%
Std. moneyness ($K/S_t - 1$)	0.32%	0.48%
Average implied volatility	12.90%	16.36%
Std. dev. of implied volatility	2.79%	2.76%
Minimum implied volatility	9.04%	11.69%
Maximum implied volatility	30.96%	34.39%

Table 1 summarizes the characteristics of the one-month S&P500 index option data used in this study. The source of this data is the Berkeley Options Database (1991-1995). Each month, a cross-section of options (both puts and calls) with approximately one-month until expiration is extracted from the database. Options with moneyness between 10% and -10% are included in the sample. Data screening criteria are described in the text of the paper.

The database of option closing prices is constructed using the following procedure. First, "end-of-day" option prices are collected each day for all contracts by averaging the last recorded bid-ask quote of the day between 2:00 and 3:00 PM Central Time. The cross-section of midquotes from the last hour of trading is not entirely synchronized, since the S&P500 index level may change over the last hour of trading. To correct for this, a Black-Scholes (1973) implied volatility is calculated each day for each option contract using the end-of-day midquote, the contemporaneous S&P500 index level, the appropriate riskless interest rate, time until expiration, and dividend yield. The closing price for each contract is obtained by evaluating the Black-Scholes formula using the same inputs and the implied volatility, except that the closing S&P500 index level is used instead of the synchronized S&P500 level. Then, each call (put) price is averaged with the synthetic put (call) price obtained using put-call parity adjusted for dividends to obtain the closing call (put) price used in the estimation procedure.

The second panel of Table 1 reports characteristics of the ATM and OTM put options used in the hedging tests. The hedging sample is composed of options with approximately one-month (from 16 to 24 days) until expiration. On each sample date, the "ATM put" is the closest-to-the-money put (within 1% of the money). The OTM put is the put with moneyness closest to -3%, but no greater than -3%. When suitable options with reported prices on the sample date and next trading date cannot be found, the sample date is excluded from the analysis. The option time-until-expiration is measured in trading days.

Table 2 - Dynamic state probability models

S&P500 daily log-return summary statistics, 1970-1995

Number of observations	Annualized mean	Annualized std. dev.	Skewness	Excess kurtosis	Normality test p-value	Serial correlation test p-value	ARCH test p-value
6571	7.55%	14.79%	-2.31	60.20	<.0001	<.0001	<.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	0.0004	3.09	0.0004	4.10	0.0003	2.41
θ	0.1483	6.71	0.1363	10.62	0.1361	10.60
ω_1	7.06E-05	17.85	1.22E-06	3.45	1.26E-06	3.80
ω_2	-3.41E-05	-8.52	-4.64E-07	-1.77	-4.24E-07	-1.69
α	0.1964	2.86	0.0663	2.41	0.0262	3.12
β			0.9212	35.88	0.9264	53.83
γ					0.0672	2.41

Log-likelihood	21941.39	22332.09	22368.10
ARCH p-value	<.0001	0.3426	0.7004
LR test p-value		<.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
6571	0.0004	1.0000	-0.36	4.26	<.0001	0.7378	0.7004

Probabilities of extreme events

Probability	Empirical density	Standard normal
Innovation < -10 std. dev.	0.00030	<.000001
Innovation < -5 std. dev.	0.00061	<.000001
Innovation < -3 std. dev.	0.00427	0.00135
Innovation > 3 std. dev.	0.00290	0.00135
Innovation > 5 std. dev.	0.00015	<.000001
Innovation > 10 std. dev.	<.000001	<.000001

The first panel of Table 2 reports summary statistics for daily S&P500 returns over the period from 1970-1995. 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 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 sample period. Robust t-statistics are calculated according to the method of Bollerslev and Wooldridge (1992). The asymmetric GARCH model is defined in equations (13) and (14). 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 LR (likelihood ratio) test p-value measures statistical significance of the improvement of the GARCH(1,1) model versus the ARCH model and the asymmetric 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 third and fourth panel report the properties of the standardized residuals from the asymmetric GARCH model. The standardized residuals are calculated by dividing the ordinary residual (ϵ_t) by its conditional standard deviation (σ_{t-1}), 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. The fourth panel compares probabilities of extreme innovations for the empirical innovation density and a standard normal density.

Table 3 - EPK estimation results**Power specification parameter estimates**

N=53	Mean	Standard deviation	Min	Max
$\theta_{0,t}$	1.0051	0.0063	0.9866	1.0185
$\theta_{1,t}$	7.36	2.58	2.36	12.55
Forecast error standard deviation	\$0.63	\$0.26	\$0.28	\$1.34

Orthogonal polynomial specification parameter estimates

N=53	Mean	Standard deviation	Min	Max
$\theta_{0,t}$	0.19	0.10	0.04	0.40
$\theta_{1,t}$	-2.25	1.06	-4.38	-0.25
$\theta_{2,t}$	-0.88	0.68	-2.52	0.19
$\theta_{3,t}$	-1.08	0.42	-1.94	-0.19
Forecast error standard deviation	\$0.09	\$0.05	\$0.03	\$0.24

Table 3 reports results from the estimation of the power and orthogonal polynomial pricing kernel specifications. These estimates are for one-month pricing kernels defined over S&P500 return states, estimated monthly over the period 1991-1995. The power specification is defined by $M_t(r_{t+1}) = \theta_{0,t}(1+r_{t+1})^{-\theta_{1,t}}$, and the orthogonal polynomial specification is defined by $M_t(r_{t+1}) = \theta_{0,t}T_0(r_{t+1})\exp[\theta_{1,t}T_1(r_{t+1}) + \theta_{2,t}T_2(r_{t+1}) + \theta_{3,t}T_3(r_{t+1})]$.

Table 4 - Analysis of EPK risk aversion

Properties of EPK risk aversion (monthly, 1991-1995)

Number of observations	Mean	Std. dev.	Min	Max	First-order autocorrelation
53	7.36	2.58	2.36	12.55	0.45

Multiple regression of monthly EPK risk-aversion on business cycle indicators

Number of observations = 46 (see notes below)

Variable	Regression coefficient	t-statistic	p-value	Univariate correlation	Correlation p-value
Intercept	3.99	3.88	0.0004		
Lagged EPK risk aversion Previous month	0.48	3.54	0.0011	0.45	0.0018
S&P500 return Percent monthly	0.04	0.29	0.7732	0.03	0.8555
Three-month Treasury yield Percent monthly change	0.68	0.37	0.7130	0.24	0.1058
Credit spread (Baa - 30 year Treasury) Percent monthly change	9.95	2.12	0.0405	0.50	0.0004
Term structure slope (30 year - 3 month) Percent monthly change	0.23	0.12	0.9035	-0.36	0.0129
Implied vol. - objective vol. spread Monthly change	0.52	3.36	0.0018	0.39	0.0069
Consumption growth Percent monthly change, annual rate	-0.04	-0.72	0.4789	-0.10	0.5224
Adjusted R-square	46.10%				

Table 4 presents an analysis of properties of EPK risk-aversion, which is estimated monthly over the period 1991 – 1995, using a power pricing kernel specification. The first panel contains summary statistics for EPK risk-aversion. The second panel reports results of a multiple regression of EPK risk-aversion on business cycle indicators. There are seven months for which insufficient data exists for estimation of the EPK, leaving 53 observations. There are seven months for which there is no one-month lagged EPK risk-aversion, leaving 46 observations. Independent variables are measured on the same date on which EPK risk-aversion is calculated. OLS t-statistics and p-values for each regression are reported. In addition, univariate correlations of the independent variables with EPK risk-aversion and p-values are reported.

The independent variables are defined as follows. The S&P500 return is the monthly percent change in the S&P500 index. The change in the implied volatility - objective volatility spread is the monthly percent change in 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). All interest rate variables are measured in percentage terms. The change in the credit spread is the monthly change in the difference between average yield-to-maturity on Moody's Baa rated seasoned long term corporate bonds and the thirty-year constant maturity treasury bill yield, using data from the Federal Reserve's H.15 release. The change in the slope of the term structure is the monthly change in the difference between the thirty-year and three-month constant maturity Treasury yield, using data from the H.15 release. The change in the three-month T-bill yield is the monthly difference between the three month constant maturity Treasury yield reported in the H.15 release. Consumption growth is the percentage growth rate of per-capita non-durable goods and services (monthly, real, seasonally adjusted) from the Federal Reserve's FRED database. Monthly U.S. resident population estimates are from the Census Bureau.

Table 5 - Models to forecast the one-day ahead pricing kernel

Time-invariant power specification

N=53	Estimate
$\theta_{0,t}$	1.0051
$\theta_{1,t}$	7.3629

EPK power specification

					Average one-day parameter forecast	Std. dev. one-day parameter forecasts
N=133	α	β_0	β_1	Adj-R ²		
$\theta_{0,t}$	0.8163	0.1885	0.0000	12.25%	1.0053	0.0030
p-value	<.0001	<.0001	0.8632			
$\theta_{1,t}$	2.2608		0.6947	49.92%	7.4872	1.8717
p-value	<.0001		<.0001			

Time-invariant orthogonal polynomial specification

N=53	Estimate
$\theta_{0,t}$	0.1945
$\theta_{1,t}$	-2.2457
$\theta_{2,t}$	-0.8791
$\theta_{3,t}$	-1.0752

EPK orthogonal polynomial specification

							Average one-day parameter forecast	Std. dev. one-day parameter forecasts
N=133	α	β_0	β_1	β_2	β_3	Adj-R ²		
$\theta_{0,t}$	0.1756	0.4321	0.0290	-0.0075	-0.0065	39.45%	0.2210	0.0698
p-value	0.0739	0.0996	0.1568	0.8530	0.8961			
$\theta_{1,t}$	0.1845		0.8545	-0.4116	0.2034	51.72%	-2.0464	0.7858
p-value	0.0429		<.0001	0.0556	0.6154			
$\theta_{2,t}$	-0.1454		0.1384	0.3496	0.0240	33.60%	-0.7196	0.4056
p-value	0.2665		0.2823	0.0217	0.9330			
$\theta_{3,t}$	-0.2535		0.1455	-0.1492	0.5352	45.67%	-0.9768	0.3011
p-value	0.0024		0.0729	0.1166	0.0033			

Table 5 reports estimates of the models used to forecast the one-day ahead pricing kernel for use in the hedging tests. The time-invariant power specification model is obtained by averaging the scaling factor ($\theta_{0,t}$) and the risk-aversion coefficient ($\theta_{1,t}$) using the 53 monthly estimates from the EPK estimation. The time-invariant orthogonal polynomial specification model is obtained by averaging the scaling factor ($\theta_{0,t}$) and generalized Chebyshev coefficients ($\theta_{1,t}$, $\theta_{2,t}$, $\theta_{3,t}$) using the 53 monthly estimates from the EPK estimation.

The EPK forecasting models are estimated using a linear regression of pricing kernel parameter estimates on one-day lagged parameter estimates and the one-day S&P500 return. For estimation, we construct a sample of estimated pricing kernels (power and orthogonal polynomial) on each of the 243 hedging sample dates. The hedging dates that are not one day apart are dropped from the sample. The remaining 133 observations are used in the forecasting model estimation regressions. The scaling factor ($\theta_{0,t}$) is only included in its own forecasting regression.

The coefficients on the lagged parameter values (α , β_0 , β_1 , ...) are reported along with their p-values and each regression adjusted r-squared. Parameter forecasts are obtained on each of the 243 hedging dates using the estimated models. Their average values and standard deviations are also reported.


Table 6 - Hedging test results

	Standard deviation of daily price changes	Reduction in hedging error standard deviation (EPK versus time-invariant specification)	Robust t-statistic (performance of EPK versus time-invariant specification)	Average number of units of underlying per written OTM put	Average number of units of ATM option per written OTM put
Portfolios - pricing kernel specification					
No hedge:					
\$100 OTM written put position	\$22.56			0.000	0.000
Hedge using underlying:					
Time-invariant power pricing kernel	\$12.41			-0.105	0.000
EPK power pricing kernel	\$12.11	2.39%	1.16	-0.100	0.000
Time-invariant orthogonal polynomial pricing kernel	\$13.45			-0.190	0.000
EPK orthogonal polynomial pricing kernel	\$13.13	2.36%	1.95	-0.116	0.000
Hedge using ATM put:					
Time-invariant power pricing kernel	\$11.21			0.000	0.298
EPK power pricing kernel	\$11.10	0.95%	2.82	0.000	0.293
Time-invariant orthogonal polynomial pricing kernel	\$11.99			0.000	0.367
EPK orthogonal polynomial pricing kernel	\$11.64	2.90%	1.74	0.000	0.256
Hedge using underlying and ATM put:					
Time-invariant power pricing kernel	\$11.36			-0.024	0.231
EPK power pricing kernel	\$11.29	0.63%	2.94	-0.022	0.229
Time-invariant orthogonal polynomial pricing kernel	\$12.08			0.083	0.526
EPK orthogonal polynomial pricing kernel	\$11.67	3.39%	2.12	0.080	0.428

Table 6 reports the hedging test results comparing the EPK power and orthogonal polynomial pricing kernel specifications with the time-invariant (average) power and orthogonal polynomial specifications over the period 1991 – 1995 (243 sample dates). The particular hedging problem chosen is hedging a \$100 position in one-month out-of-the-money (OTM) S&P500 index put options using one-month 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 using the methodology described in the paper.

The second column reports the daily price change standard deviation and may also be interpreted as a daily percentage hedging error, since the initial portfolio value is \$100. The third column reports the percentage reduction in hedging error obtained by switching from the time-invariant pricing kernel to the time-varying pricing kernel. The fourth column reports a test of the statistical significance of the improvement in hedging performance in hedging using a time-varying pricing kernel versus a time-invariant pricing kernel. This test is implemented using a predictive accuracy statistic similar to that proposed by Diebold and Mariano (1995). The relative performance improvement is measured using the t-statistic for the squared error loss, which is defined as the average difference between the squared time-invariant pricing kernel one-day hedging error and the squared EPK pricing kernel one-day hedging error. Standard errors are calculated using the heteroskedasticity and autocorrelation consistent covariance matrix of Newey and West (1987). When the t-statistic is greater than 2 and positive, the EPK pricing kernel hedge performance is significantly better than the time-invariant pricing kernel hedge performance. The average hedge portfolio weights are reported in the last two columns of the table.

Figure 1
One-month S&P500 index option prices



This figure graphs the cross-section of one-month S&P500 option prices from June 1991 – 1995 against dollar moneyness (option exercise price – closing S&P500 index level). For negative dollar moneyness, put premia are reported. For positive dollar moneyness, call premia are reported. Hence, all option prices depicted are for out-of-the-money contracts.

Figure 2
Estimated S&P500 return density functions
June 1991, June 1992, ... , June 1995

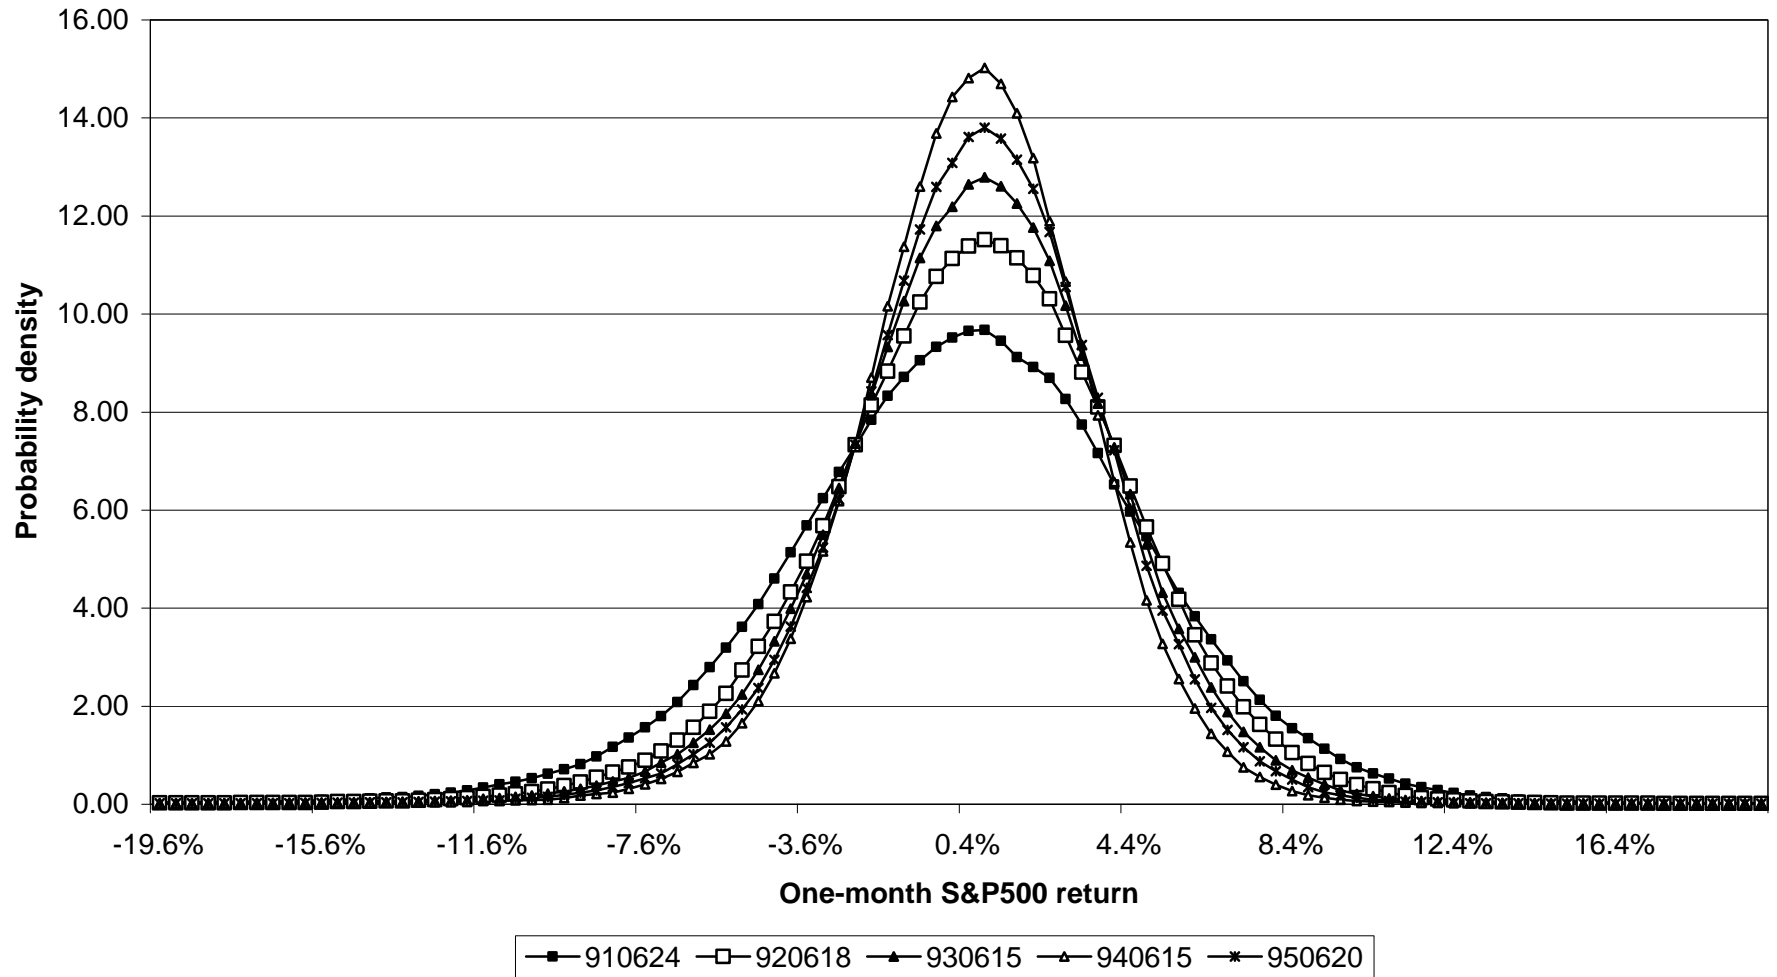


Figure 2 graphs the estimated S&P500 return probabilities density functions for five representative months (June 1991, June 1992, ..., June 1995). This density function is obtained using Monte-Carlo simulation (200,000 replications) of the estimated asymmetric GARCH model with an empirical innovation density. The plotted density functions are obtained using a Gaussian kernel density estimator with bandwidth equal to $.9 \cdot N^{-1/5} \sigma$, where σ is the return standard deviation and N is the number of replications.

Figure 3
Annualized S&P500 conditional volatility
Asymmetric GARCH model (1970-1995)




Figure 4 graphs the annualized conditional volatility forecasts for the asymmetric GARCH model over the period 1970-1995.

Figure 4
Empirical pricing kernel: power specification
June 1991, June 1992, ..., June 1995




Figure 4 graphs five representative one-month empirical pricing kernels for June 1991, ... , June 1995 using the power pricing kernel specification: $M_t(r_{t+1}) = \theta_{0,t}(1+r_{t+1})^{-\theta_{1,t}}$. The time variation in the one-month empirical pricing kernel suggests that constant relative risk aversion over equity return states is not supported by the data.

Figure 5
Empirical pricing kernel: Orthogonal polynomial specification
June 1991, June 1992, ..., June 1995




Figure 5 graphs five representative one-month empirical pricing kernels for June 1991, ... , June 1995 using the orthogonal (generalized Chebyshev) polynomial pricing kernel specification with four terms: $M_t(r_{t+1}) = \theta_{0,t}T_0(r_{t+1})\exp[\theta_{1,t}T_1(r_{t+1}) + \theta_{2,t}T_2(r_{t+1}) + \theta_{3,t}T_3(r_{t+1})]$. This specification is more flexible than the power specification, and the data appears to assign greater weight to large negative S&P500 return states and lesser weight to large positive S&P500 return states.

Figure 6
Time-invariant (average) pricing kernels
June 1991, June 1992, ..., June 1995

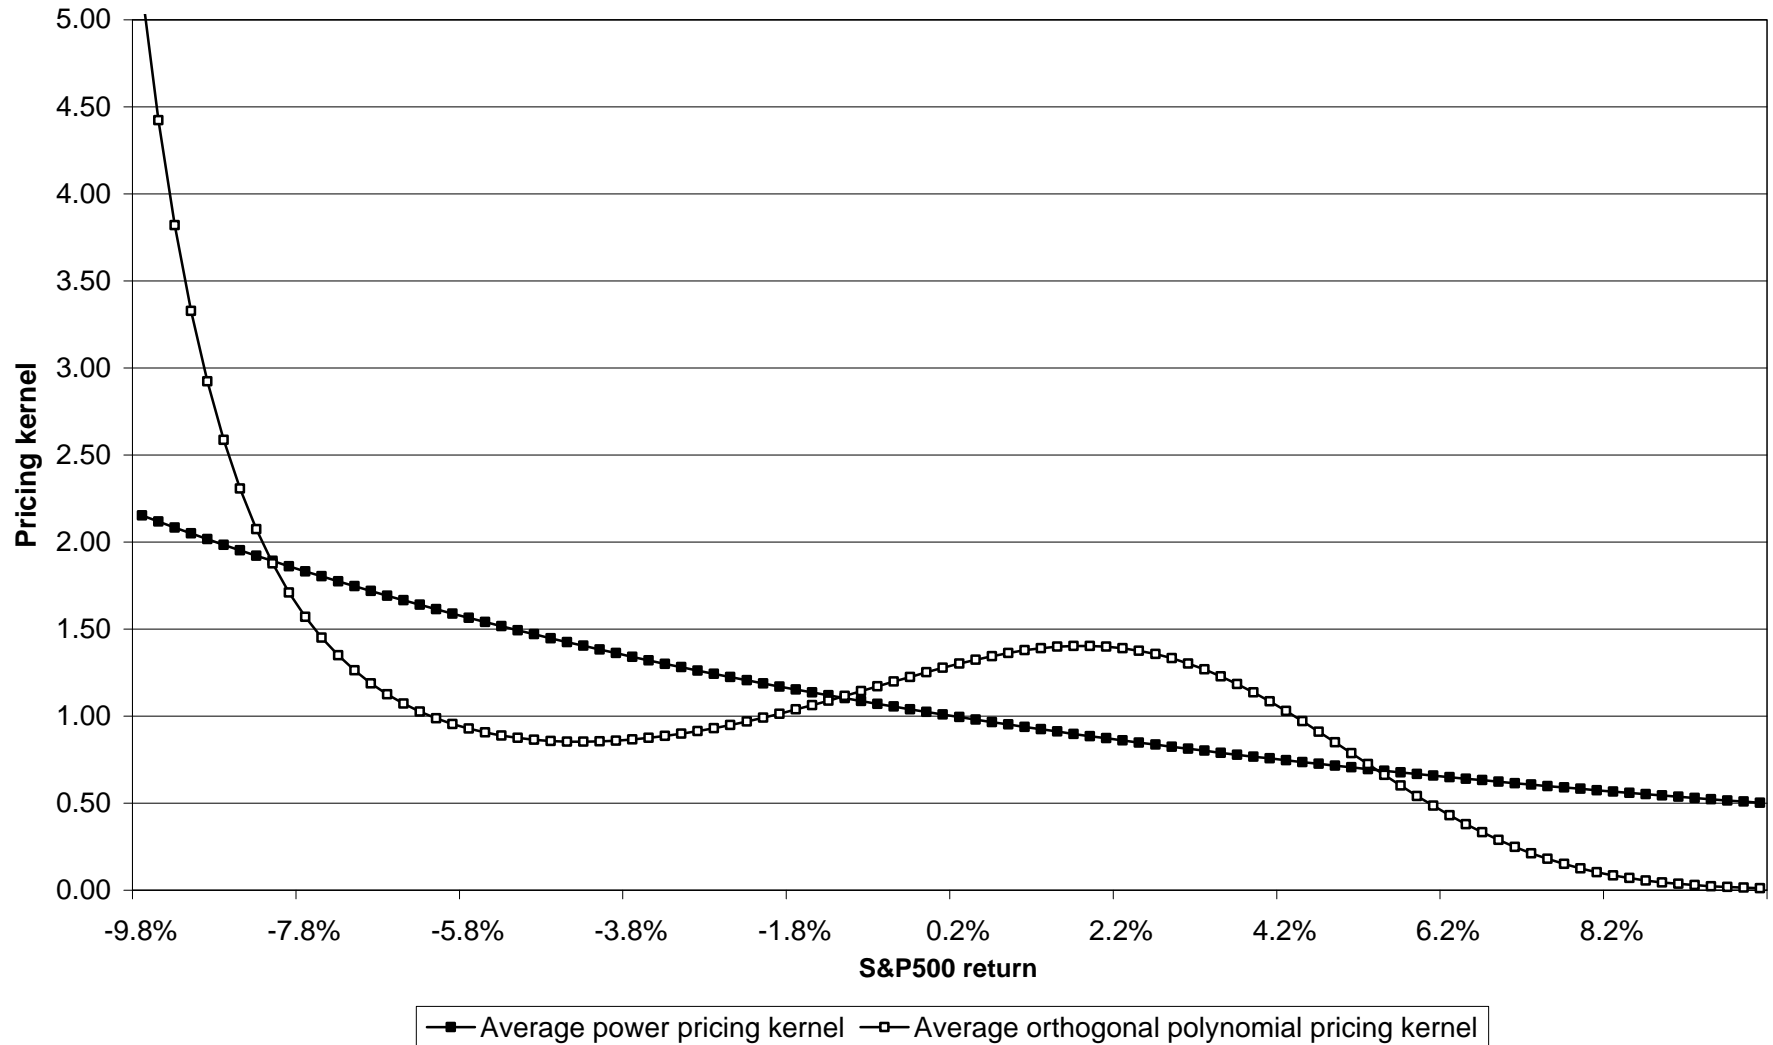


Figure 6 graphs two time-invariant (average) pricing kernels using the power and orthogonal polynomial pricing kernel specifications. Each average pricing kernel is obtained by evaluating the pricing kernel specification at the average parameter estimates over the period from 1991 to 1995.

Figure 7
Monthly EPK risk aversion




Figure 7 graphs EPK risk aversion each month from 1991 through 1995. EPK risk aversion is the exponent of the estimated power pricing kernel ($\theta_{1,t}$). The seven missing monthly observations are indicated by absence of a data point as well as a gap in the connecting curve.