The Pricing Kernel Puzzle:
Reconciling Index Option Data and Economic Theory

by

David P. Brown

and

Jens Carsten Jackwerth*

First draft: May 9, 2000

This version: February 7, 2004

C:\Research\Paper18\paper29.doc

* David Brown is from the University of Wisconsin at Madison, Finance Department, School of Business, 975 University Avenue, Madison, WI 53706, dbrown@bus.wisc.edu Tel: 608-265-5281, Fax: 608-265-4195. Jens Jackwerth is from the University of Wisconsin at Madison and from the University of Konstanz, PO Box D-134, 78457 Konstanz, Germany, jens.jackwerth@uni-konstanz.de, Tel. +49-(0)7531-882196, Fax: +49-(0)7531-883120. The authors thank Jim Hodder, Stewart Hodges, Lars Norden, Mark Rubinstein, and seminar participants at CEPR Summer Symposium at Gerzensee, Tilburg University, European Investment Review Conference - Paris, Insead, London School of Economics, Warwick University, Lancaster University, University of Lausanne, 2nd World Congress of the Bachelier Society - Heraklion for valuable comments. The authors gratefully acknowledge a research grant from INQUIRE UK. This article represents the views of the authors and not of INQUIRE.

The Pricing Kernel Puzzle:

Reconciling Index Option Data and Economic Theory

Abstract

The pricing kernel puzzle of Jackwerth (2000) concerns the fact that the empirical pricing kernel implied in S&P 500 index options and index returns is not monotonically decreasing in wealth as standard economic theory would suggest. Thus, those options are currently priced in a way such that any risk-averse investor would increase his utility by trading in them. We discuss the limited literature, which deals with the puzzle directly, but find that those papers cannot explain the puzzle, either. We provide a novel representative agent model where volatility is a function of a second momentum state variable. This model is capable of generating the empirical patterns in the pricing kernel, albeit only for parameter constellations which are not typically observed in the real world.

The Pricing Kernel Puzzle:

Reconciling Index Option Data and Economic Theory

The Capital Asset Pricing Model (CAPM) of William Sharpe and the option pricing models of Fisher Black, Robert Merton, and Myron Scholes were seminal in developing our understanding of the pricing of financial assets; these works sparked a firestorm of research by economic theorists and empiricists. Despite the fact that the CAPM was developed to price equity shares while the option pricing models were developed for options, these two sets of models share a common feature, namely a state-price density, which allows us to price all securities. Early researchers understood the original works as distinct theories. The authors of the option models employed different assumptions and logic than Sharpe, and explained their results using a disparate set of arguments. However, it is now understood that asset-pricing models generally – a set that includes the CAPM, option pricing models, and the work of Debreu (1993), Rubinstein (1976), and Cox, Ingersoll and Ross (1985) – can be interpreted using a single, simple calculation. An asset price is equivalent to an expectation based on a state-price density (SPD).[1]

Because of its central role in pricing assets, the state-price density now has many uses. For example, Campa and Chang (1998) describe how central banks use estimates of SPDs to measure the credibility of exchange rate commitments and Jackwerth (1999) lists numerous other applications. Further use for SPDs is in risk management by banks, trading firms, and portfolio managers. Because an appropriate hedge position - say a position in an option - requires an assessment of the random changes in the value of the option, risk managers often rely on pricing models to evaluate the hedge. In turn, accurate valuation of the hedge requires precise estimates of the SPD.

Not surprisingly, researchers have taken great interest in estimating SPDs. An often-used approach relies on a model of a representative agent, and requires an estimate of the parameters of the agent's utility function. This is typically accomplished by fitting the optimality conditions

---

[1] For simplicity and without loss of generality, we assume in our discussion that the rate of interest on risk-free assets is zero. As a result, there is no discounting of the expectation needed. In general, the price of a security is the sum of expected future payoffs (using the SPD) discounted at the risk-free rate. The word "density" is usually reserved for a continuous function. For simplicity of discussion, we still refer to the discrete risk-neutral probability distribution under which the expectation is taken as a state-price density.

of the agent's utility maximization to a set of asset returns and other data, e.g., per capita consumption. Two of many examples are Brown and Gibbons (1985) and Hansen and Singleton (1983). An alternative method is the subject of this research. This method draws on a cross section of options with different strike prices and does not impose strong assumptions on the utility of a representative agent. In fact, this method does not even require that a representative agent exists.

Breeden and Litzenberger (1978) first showed within a theoretical setting, given a set of options on aggregate consumption dense in the set of possible strike prices, that a state-price density could be calculated exactly from the option prices. As Cox, Ingersoll, and Ross (1985) demonstrate, this SPD may be used to price all securities. Rubinstein (1994) and Jackwerth and Rubinstein (1996) provide empirical procedures for estimating a SPD when a finite number of options exist instead of a dense set. They also extend the interpretations of Breeden and Litzenberger to the case where options are written on securities or security indices, and not on aggregate consumption. In this case, the SPD estimate is a projection; for example, the SPD obtained from options on an equity index is an empirical projection of the general consumption-based SPD (discussed by Cox, Ingersoll, and Ross (1985)) onto the space of payoffs defined by the index. This implies that the SPD estimated from index option prices cannot be used to price all securities, but it can be used to price all securities in a large and important subset, namely the set of options and other securities derived from the index payoffs.

There is now a large literature documenting estimates of SPDs derived from various options markets, e.g., foreign exchange, interest rates, equities, and equity indices. Some 100 papers on this topic are surveyed in Jackwerth (1999). As the evidence regarding the state-price densities has been collected, a puzzle has arisen and was first documented in Jackwerth (2000). Estimates of the SPD derived from index options, which are written on the S&P 500 index and comprise the second largest options market in the world, appear to be inconsistent with fundamental assumptions about investor behavior. Specifically, when recent estimates are viewed through the lens of existing asset pricing theory, they imply that investors behave as if they are risk seeking (and not risk averse) with respect to the risk inherent in the index. This conclusion is inconsistent with one of the fundamental assumptions of the economic theory. The empirical research suggests that the marginal utility of investors trading S&P 500 index

portfolios is increasing over an important range of wealth levels and not decreasing in wealth, as economic theory would suggest. We seek to understand this puzzle.

There are several reasons to investigate this puzzle, and to investigate more generally the state-price densities for equity indices. A first is that indices such as the S&P 500, DAX, FTSE, Nikkei, and those published by the Center for Research in Security Prices of the University of Chicago (CRSP) represent the large majority of the public equity capital in their respective nations. The S&P 500 index represents roughly 50% of public U.S. equity capital.

Second, the very large market capitalization of broad equity indices has led to a long line of research on their distributional properties. For example, research into the random walk theory of equity price changes has asked if the S&P index returns are predictable. Also, early tests of the CAPM were carried out under the assumption that the value-weighted CRSP index was a good proxy for the entire capital of the US investors. More recent works, for example by Fama and French (1988) and Poterba and Summers (1988), examine the predictability of a number of stock market indices. By their very nature, state-price densities provide information regarding the statistical distribution of index returns.

In the following section 1 we describe the pricing kernel puzzle in more detail. In section 2 we discuss the literature and find that not many papers address the pricing kernel puzzle at all. Moreover, the few that do also have some problems which we detail. We therefore suggest in section 3 a novel model based on a representative investor who is faced with a second momentum state variable. Section 4 concludes. Information concerning the data is contained in the appendix.

Section 1 - Fundamental Determinants of State Prices

We introduce the SPD by calculating the price of a European option. The option price C is the expectation of the random payoffs received at the option's expiration. Write C as a sum across states,

$$C = \sum_{s} q_s F_s,\qquad(1)$$

where F_s is the option payoff and q_s is the current Arrow-Debreu state price of a dollar paid in the future state s.[2,3] The sum of the q_s across states is equal to one, so the set of q_s represent a distribution analogous to, but distinct from a market participant's subjective distribution of probabilities of the states. Equilibrium theory relates a state-price density to economic fundamentals. Breeden and Litzenberger (1978) and Cox, Ingersoll, and Ross (1985) show that a state price q_s satisfies

$$q_s / p_s = m_s \qquad (2)$$

where p_s is the subjective probability assessed by an investor for state s, and m_s is known as the pricing kernel and is proportional to the marginal utility of a representative investor in that state. In a state where the marginal value of an additional dollar is high, the proportional difference between the state price and the subjective probability is high. Conversely, if one observes a state price far below a subjective probability, then the marginal value of a dollar is low in that state.

Within an exchange economy model such as Rubinstein (1976) and Lucas (1978), the pricing kernel is the ratio of marginal utilities of a representative investor. In this case the value of an option written on the market portfolio is

$$C = \sum_s m_s p_s F_s = \sum_s \left(U'(C_s)/U'(C_0)\right) p_s F_s, \qquad (3)$$

where m is a ratio of marginal utilities of future per capita consumptions (C_s) and present per capita consumption (C_0). This insight – that the pricing kernel is proportional to a ratio of marginal utilities - has led numerous authors to investigate the distributions of equity returns and test a variety of alternative specifications of utilities and endowments for the representative agent. One of the first utility functions tried was power utility. Described in detail by Rubinstein (1976) and investigated empirically by Brown and Gibbons (1985) and others, an

---

[2] See Debreu (1993) for development of a state price.
[3] We focus our attention on European options because they are exercised only at expiration. A European call option pays the maximum of zero and the difference between the price of the underlying security, e.g. the level of a stock index P_s, and a strike price X defined in the option contract. We write F_s=Max[0,P_s-X].

implication of power utility is that aggregate consumption is proportional to aggregate wealth.[4] In this case, the pricing kernel m is a nonlinear function of the return on the index measured over the life of the option. Furthermore, because the power-utility investor is risk averse, m is a monotonically decreasing function of aggregate wealth at the option expiration.

Jackwerth (2000) estimates q and p using data on the S&P 500 index return and option prices from April 2, 1986 to the end of 1995. Details pertaining to the data are in the appendix. Estimates of m are then obtained from equation (2). The maximum smoothness method of Jackwerth and Rubinstein (1996) is applied to the option data and an estimate of q is obtained. This method does not require a representative agent to be specified. Multiple options with various strike prices, but all with 31-day lives are used. Post crash option prices exhibit so-called implied volatility smiles and skews.[5] For example, researchers often describe the cross section of options by a plot of implied volatilities against moneyness (= ratio of strike prices to index level) as in Figure 1.

---

[4] This result also requires that per capita consumption follows a geometric random walk with drift.

[5] That is, the volatility (of returns on the underlying asset) consistent with the Black-Scholes model is obtained, which exactly matches the observed option price. This is the "implied volatility," which is used as a substitute for the option price and makes it an easy task to compare the prices of options.

Figure 1. Empirical Volatility Smile. Typical post-1987 stock market crash implied volatility smile across moneyness (= ratio of strike prices to index level).
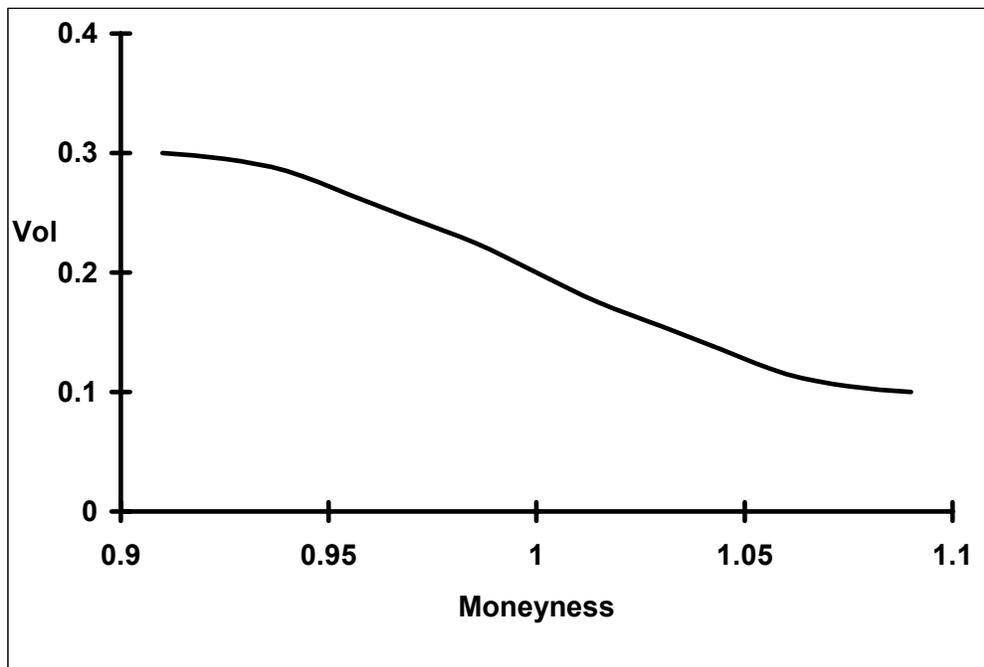

A consequence of the Black-Scholes option-pricing model is that implied volatilities of all options are equal. If this result held, the plot in Figure 1 would be flat. Instead, we observe in post crash index options a "smile", or "skew", or "smirk" depending on your preference in label. The left tail of the plot is substantially above the right.[6] An implication is that the risk-neutral distribution q is left-skewed and leptokurtic, that is, it has a fat left tail and is more peaked in the center than the lognormal distribution.

The subjective probability distribution p is estimated as the historical distribution of four years worth of index returns over 31-day intervals. This is an appropriate method for estimating p provided that investors have rational expectations.[7] Typical estimates of the distributions, say q* and p*, are graphed in Figure 2.

---

[6] For example, for options on individual equities, the plot is relatively flat compared to Figure 1. For foreign-exchange options, each tail of the plot tends to be slightly and equally elevated.

[7] Investors also need to be solely concerned with the unconditional distribution, or they need to assume that the distribution of returns is stationary. Alas, work by Rosenberg and Engle (2002) uses a non-stationary GARCH model but finds similar subjective distributions.

Figure 2. Risk-Neutral and Subjective Distributions. The risk-neutral and the subjective return distributions are calculated on March 16, 1990. The subjective distribution is approximated by the smoothed 4-year historical distribution. Returns are reported as 1 plus the rate of return.
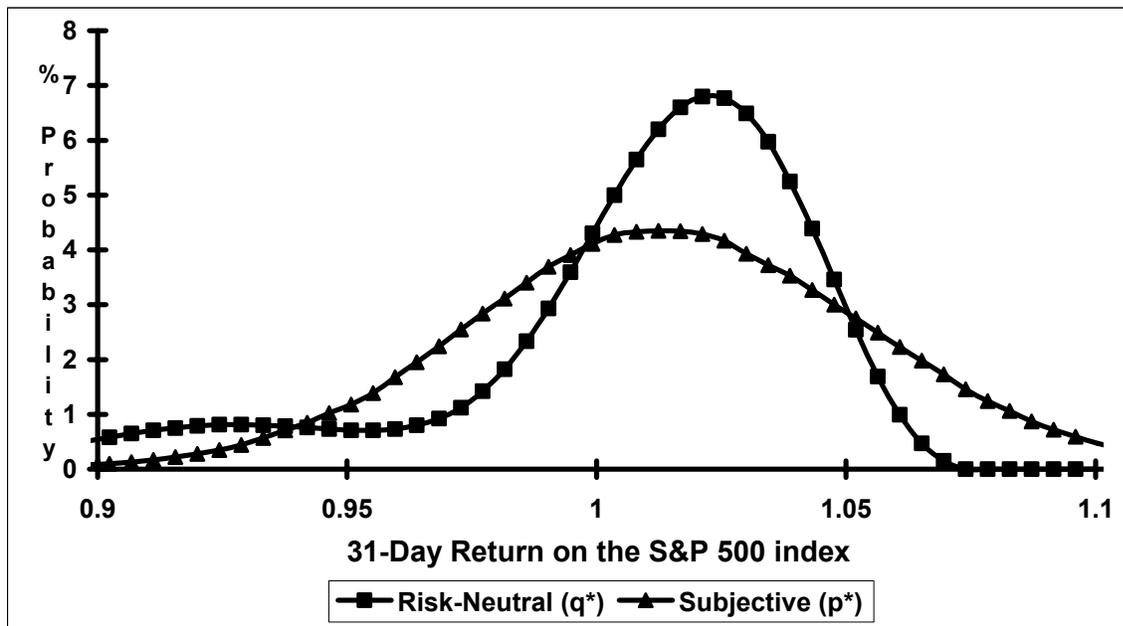
A typical estimate of m, say m*, appears in Figure 3 as a function of the return on the S&P 500 index. Rosenberg and Engle (2002) confirm this shape by using a related method, and Jackwerth (2004) provides international evidence for Germany, Japan, and the UK.[8]

---

[8] In a further exploratory study, we investigated market frictions in a GMM setting starting with the Euler equation E[m r] = 1. We use the dataset on options returns from Buraschi and Jackwerth (2001). He and Modest (1995) suggest that the vector of 1's should be replaced with some other value, close to 1 for transaction costs or less than 1 for short sale constraints. However, as the marginal investor (the one with the lowest costs) sets the equilibrium prices, it is hard to imagine that the vector should be vastly different from 1; especially since the market maker and the traders are not short sale constraint. We need to use an expected value of the OTM put returns, which is some 40-50 basis points lower than 1, in order to achieve a monotonically decreasing pricing kernel. Such large daily transaction cost is hard to imagine in this liquid market. Such diminished return could also suggest mispricing, a position explored in Jackwerth (2000) and supported by Agarwal and Naik (2000), who find that hedge funds tend to sell out-of-the-money and at-the-money puts. Bondarenko (2003), Coval and Shumway (2001), and Driessen and Maenhout (2003) also find that similar option trading strategies turn out to be very profitable. Given the size and liquidity of this market, our prior is however, that such profit opportunities should not exist for ten years without any tendency towards equilibrium.

Figure 3. Empirical Pricing Kernel. Typical post-1987 stock market crash implied pricing kernel. The pricing kernel is calculated as the ratio of the option implied state-price density and the historical smoothed return distribution. Returns are reported as 1 plus the rate of return.
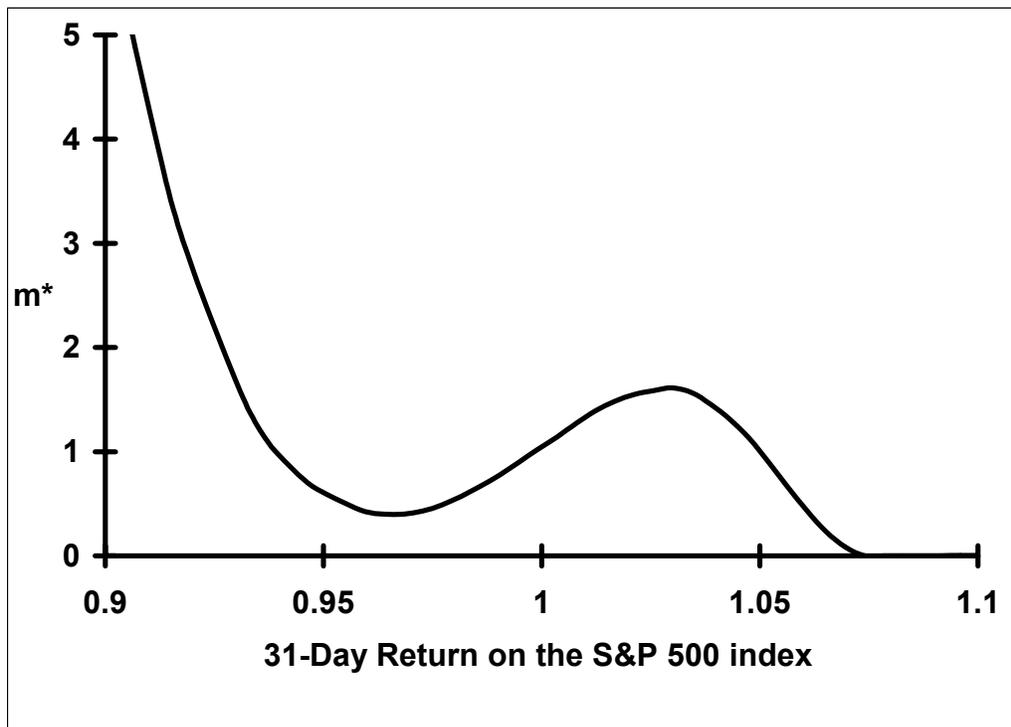
For any option-pricing model, the plot of the kernel m (as in Figure 3) is closely tied to the degree of the smile (like that of Figure 1).[9] It is important to note that any model used to explain the smile in the S&P 500 index options must provide a kernel m that is consistent with Figure 3. That is, the kernel implied by a well-constructed model must have a central range that is increasing with the index. The goal of this work is to find models that explain the pricing kernel as shown in Figure 3. The value of 1 in the center of the horizontal axis represents an ending level of the index (i.e., at the time of option expiration) that is equal to the current level (at the beginning of the 31-day interval). Globally, $m*$ is a decreasing function of the ending index level. However, for the range from approximately 0.97 to 1.03, i.e. for a range of index levels centered on and within a 3% deviation from the current level, $m*$ is increasing. This

---

[9] We do not formally demonstrate the link between plots of m and smiles here, but refer to the discussion in Jackwerth (2000) and Rosenberg and Engle (2002).

occurs because the proportional difference between the estimates, say $q*$ and $p*$, is increasing over this range, as equation (2) implies.

Traditional asset pricing theory, e.g. Rubinstein (1976) and Lucas (1978), assumes that a representative investor exists. It also is common in tests of asset pricing theories to assume that a market index such as the S&P 500 index represents the aggregate wealth held by this investor. If we make these assumptions, then the empirical $m*$ of Figure 3 suggests that the representative investor is locally risk seeking. Over the range of 0.97 to 1.03, the marginal utility is increasing in wealth, the utility function is convex, and the investor will pay to acquire fair gambles in wealth. Taking this thought experiment to the limit, however, we recognize that our interpretation of $m*$ in equations (2) and (3) is derived from the optimality conditions of an investor with concave utility. Figure 3 is inconsistent with these conditions. Hence we arrive at a puzzle. Our estimate $m*$ is inconsistent with standard approaches in asset pricing.

Our estimation technique does not require the above strong assumptions. To estimate q Jackwerth and Rubinstein (1996) only require that the prices do not offer any arbitrage trading opportunities, while the sample estimator of p requires rational expectations to be valid. Therefore we interpret $m*$ somewhat differently. It is a projection of the pricing kernel m onto the space of index returns. This projection is guaranteed to exist, independent of the existence of a representative agent for the economy as a whole. Just the same, $m*$ may be interpreted as being proportional to the marginal utility of an investor who trades in the S&P 500 index and the options written on the index. In practical terms, $m*$ represents investors in an index fund that mimics the S&P 500. It is surprising to see that index-fund investors have marginal utilities that are locally increasing in the level of the index.

Before we start modeling this behavior, we first turn to the question if the pricing kernel puzzle might be just spurious. Jackwerth (2000) discusses concerns about the quality of the estimate $m*$ in detail and has ruled out methodological or data errors. Using simulated data for which we know the true q, a number of different methods generate fairly accurate estimates. Similarly, a variety of standard procedures, both parametric and non-parametric, are used to estimate the historical return distribution $p*$ and they provide equivalent results. Furthermore, there is no good reason to believe that our index return data is poor. The hump-shape of figure 3 is also statistically significant since all estimates of $m*$ in Jackwerth (2000) exhibit it.

A potential further explanation for the existence of the smile is the illiquidity of out-of-the-money options. Options traded in a market with transaction costs or other frictions are priced within an interval determined by these costs. However, the liquidity in the index option market increased ten-fold from an average daily notional volume of $250 million in 1986 to $2250 million in 1995 (Jackwerth and Rubinstein (2001)), while the steepness of the smile did not change after the 1987 crash. In addition, preliminary research by the authors finds that large trades of out-of-the-money puts have no obvious impact on the steepness of the smile.[10]

Further, while the bid/ask spreads in relation to the option prices are bigger for away-from-the-money options, the hump in the pricing kernel occurs around at-the-money, and is determined by the at-the-money options. Hentschel (2002) suffers from the same criticism when he suggests that option quotes are distributed symmetrically around some true value and are truncated at the intrinsic value. Again, the pricing kernel puzzle arises in the center whereas the truncation mainly affects the tails of the distribution. Also, it is not clear that market participants do indeed quote options in this particular way.

Market frictions due to margin account requirements might also play a role here. The public, the clearing firms, and the market makers are all subject to margin requirements. However, the requirements have not changed dramatically around the crash of 1987. They were slightly tightened in June of 1988 and more or less remained at that level since then until the end of the sample. To explain the steep smile with transactions costs for trading the out-of-the-money options is probably insufficient, too. For one, transaction costs existed in the precrash period as well as the postcrash period without changes on the day of the crash. However, the smile changed dramatically. Second, transaction costs should affect the out-of-the-money options on both sides of the smile equally. Instead, the index option smile is largely asymmetric.

Although transaction costs are not a likely reason for the index smile, a related concern is the hedging of out-of-the-money puts. Hedging is difficult for sold out-of-the-money puts when one uses only the underlying to hedge. The hedge must offset the positive delta of the position, which is low and grows large as the market crashes. The hedging strategy must have a negative delta that declines to minus one as the market crashes. Thus, the investor must sell into the

---

[10] However, Dennis and Mayhew (2000) test for the impact of "market sentiment" variables on the smiles of individual stock options. They find that the ratio of put to call volume is significant and that higher put volume increases the slope of the smile. Still, other measures of "market sentiment", such as the Consumer Confidence Index or the average Price/Earnings ratio for the S&P 500 stock, are not significant in explaining the slope of the smile.

falling market. However, this problem only applies to investors who desire to stay fully hedged. Investors similar to the representative investor should be able to follow sell and hold strategies even without hedging since we are assuming the existence of a pricing kernel in the economy. A static unhedged position is then still utility-improving.

Note that the historical sample estimate $p*$ is backward looking. On the other hand, $q*$ is derived from option prices, and is forward looking in that it represents the beliefs of investors over the life of the index options (31 days in our work). If investors do not have rational expectations, then $q*$ and $p*$ may differ greatly. One form of irrationality we call crash-o-phobia or a peso problem. To understand this, suppose that in the period of the 1987 crash investors suddenly became overly fearful of a future market crash. Thus in the post-crash period prices of options with strike prices below the current index level are elevated relative to rational levels. As a result, the left tail of $q*$, which represents prices for dollars received conditional on a large market decline, is raised upward. Given that $q*$ always integrates to one, $q*$ needs to decline in other ranges. If in addition the statistical distribution $p*$ is unchanged – so that the change in $q*$ represents an irrational change in beliefs - then the pricing kernel may take on the shape shown in Figure 3.

Just the same, it is odd that $m*$ is locally increasing in all months of our post-crash index sample that currently ends in 1995 when the CBOE stopped distribution of the data. The implication is that crash-o-phobia continued for at least eight years after 1987, and it did so in the absence of another crash. We expect that if investor beliefs changed irrationally in the 1987 crash, then the shape of $m*$ should revert to its precrash form as memories fade during the postcrash era.[11] Moreover, the S&P 500 index options trade in a large and liquid market, and the crash-o-phobia hypothesis implies that these options are irrationally priced. If so, then over time traders should learn to profitably exploit such mispricing. Over an eight-year period arbitrage trading should drive option prices and the pricing kernel back to their rational precrash levels. However, we observe no such adjustments. For this reason, we seek rational alternatives to the crash-o-phobia hypothesis as an explanation of the puzzling pricing kernel.

---

[11] See Driessen and Maenhout (2003) for such model with irrational investors.

Section 2 - Alternative Solutions to the Pricing Kernel Puzzle

The extant literature is largely unsatisfactory in explaining the puzzle of Figure 3. Many authors model the smile that appears in Figure 1, but this literature generally is one-sided and does not tackle the puzzle. Either the index return process in modeled (generating a p distribution) or post-crash option prices are modeled (typically generating a left-skewed q distribution), but a model of m = q/p is not offered.

Consider the leverage effect as an example. As the market value falls, the market debt-equity ratio and therefore index volatility rises. Hence we should expect to see the smile in implied volatilities, as shown in Figure 1. Note that Black (1976), Christie (1982), and Schwert (1989) discuss the leverage effect on the statistical distribution, and similarly Toft and Prucyk (1997) and Dennis and Mayhew (2000) discuss the effect on the risk-neutral distribution. However, none of these papers calculates the pricing kernel or links the two distributions.

Other examples are the correlation effect (e.g., Kelly (1994), Campbell et al (2001)) and the learning or information aggregation effect (e.g., Grossman (1988), Gennotte and Leland (1990), Jacklin, Kleidon, and Pfleiderer (1992), Kleidon (1994), Romer (1993), Hong and Stein (1999) and Nandi (1999)). The correlation effect says that in down markets correlations across companies increase so that the impact of bad news on the market index is asymmetric. The learning (or information) effect is that an investor learns through trading, and occasionally prices adjust sharply despite the absence of a news event. Each of these effects can imply that volatility increases as the index declines, and thus imply that a smile exists as it appears in Figure 1. Again, none of the papers listed here model p and q jointly or describe m.

Papers that model the pricing kernel directly include Franke, Stapleton, and Subrahmanyam (1998 and 1999), who suggest a pricing kernel with declining elasticity, which follows from undiversifiable background risk faced by investors. Alternatively, Benninga and Mayshar (1997) consider heterogeneous investors, while Campbell and Cochrane (2000) explore utility functions with habit persistence. Grossman and Zhou (1996) describe an economy with demanders and suppliers of portfolio insurance and solve for m. In each of these cases, marginal utility is monotonic in wealth and m is downwards sloping.

The work of Bartunek and Chowdhury (1997) and Bliss and Panigirtzoglou (2004) suffers from a similar problem when they fit the risk aversion coefficient of a power utility

investor in order to explain observed returns best. For that they construct a subjective distribution based on the option implied risk-neutral distribution which they transform through the pricing kernel. But as they assume a power utility, their pricing kernel will by construction be monotonically decreasing. They cannot find a pricing kernel puzzle, even if it is in the data.

A large literature investigates general stochastic processes for the market index together with other variables to explain the smile in implied volatility. Some of these models allow stochastic volatility, e.g. the continuous-time model of Heston (1993) and the GARCH model of Duan (1995), or stochastic interest rates. Others consider jumps in the index or in the volatility of the index, e.g. Bakshi, Cao, Chen (1997), Bates (2000), and Pan (2002). These authors assume that jump risk is priced, but also that the compensation for jump risk (in the expected return) is proportional to the intensity of the jump component. Proportional pricing is implied by Cox, Ingersoll, and Ross (1985), who assume time-additive logarithmic utility of a representative agent. To demonstrate that this implies a monotonic m, we simulate Pan's (2002) model according to her estimated parameters. The estimated pricing kernel is monotonically decreasing as shown in Figure 4.[12]

---

[12] Ziegler (2003) demonstrates this same result.

Figure 4. Simulation of the Pricing Kernel in Pan (2002). The parameters are taken from Pan (2002), the time horizon is 30 days with 10 intra-daily discrete steps in the approximation. We use 10,000 simulation runs.
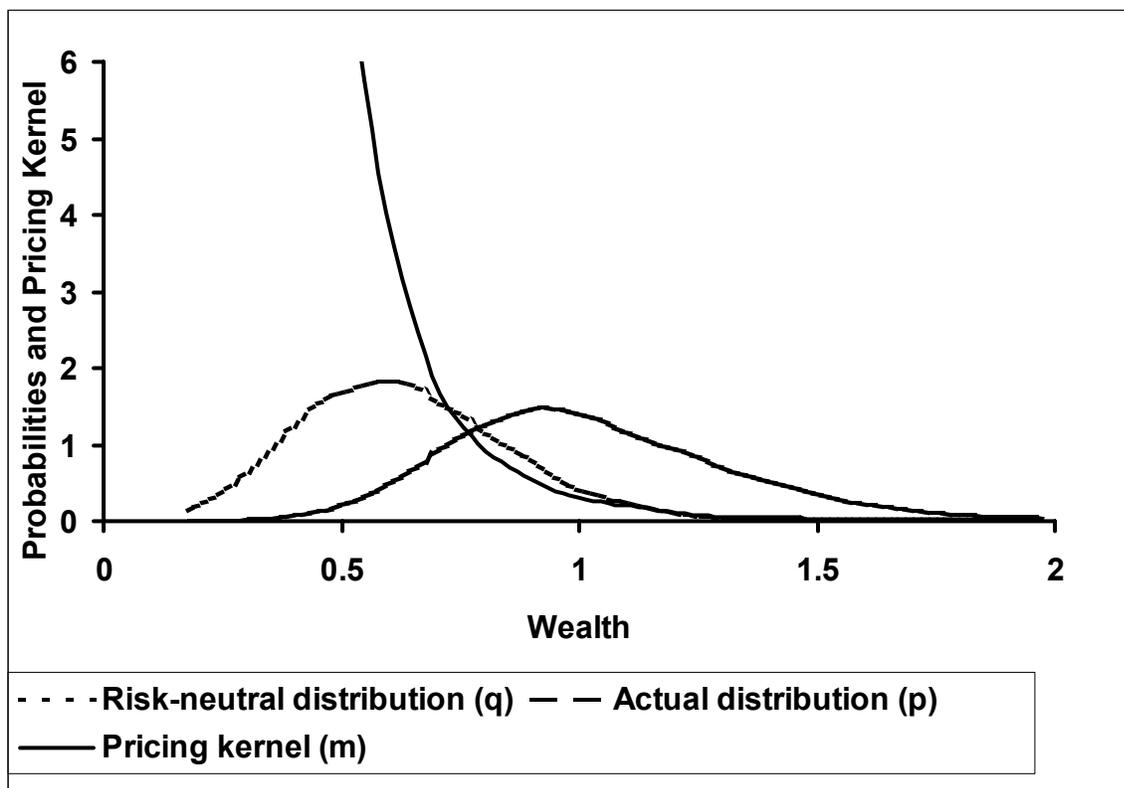
This leaves only two more papers which actually model both the q and the p distribution simultaneously: Branger and Schlag (2002) and Ziegler (2003). Branger and Schlag (2002) model an economy where the index is the average of 30 identical stocks which all are driven by jump-diffusion processes. These processes have a market component and an idiosyncratic component for both the diffusion and the jump parts, respectively. By formally changing to the risk-neutral measure, they change the drift of the diffusion but also the correlation coefficient and the intensity of the jump component. While the correlation coefficient does not have much of an impact, the jump intensity does. Ignoring the idiosyncratic components and focusing on the market component of the index, let the actual jump intensity be very small. Then, the actual return will look almost lognormal. Let the risk-neutral jump intensity be very large and the risk-neutral distribution can exhibit a pronounced left tail and a reduced mode of the distribution. The resulting pricing kernel can in such case exhibit increases in wealth.

However, Branger and Schlag (2002) are agnostic about the equilibrium implications of their mechanical change of measure. Thus, while they can generate locally increasing pricing kernels, they do not explain the underlying economics. They do not provide economic restrictions on the price of jump risk which in turn will restrict the pricing kernel in the economy.

In its intent, the work by Ziegler (2003) is closest to this paper. He considers three hypotheses: misestimation of the statistical distribution, misspecification of investor preferences, and heterogeneous investor beliefs. Using a single-factor economy with complete markets, he argues that neither misestimation nor misspecification of preferences are likely to explain the pricing-kernel puzzle. Errors in estimation of the statistical distribution, which we represent as (q* - q), must be very odd and unreasonable in their functional relation to the underlying factor (which in his model is aggregate consumption) if they are to cause the puzzle. With respect to the mispecification of preferences, he investigates commonly used preferences: constant absolute risk aversion, constant relative risk aversion, and decreasing relative risk aversion. For specifications with reasonable parameter values he finds that they all yield very similar transformation from statistical to risk-neutral distributions and vice versa.

Ziegler (2003) argues that his third hypothesis – heterogeneous beliefs – holds more promise as an explanation for the pricing-kernel puzzle. He essentially replaces the actual distribution (which is historically observable to all investors and looks rather lognormal) with a mixture of two lognormal distributions. This mixture is generated by mixing the pessimists' distribution (low mean, low volatility) and the optimists' distribution (high mean, high volatility) in proportions of roughly half and half. Such mix keeps the mean of the actual distribution, but looks rather different, namely left-skewed and leptokurtic. Still, all investors stick to their beliefs and use this mixture as the forward looking subjective distribution. Since it looks much like the risk-neutral distribution in shape, albeit shifted to the right, the resulting pricing kernel is monotonically decreasing. The model hinges on the assumption that no investor checks the statistical fit of the mixture against the historical time-series of index returns. Also, the model needs to have the polar setting of having only pessimists and optimists. A sizeable group of moderates would make the mixture look much more like a lognormal distribution and the pricing kernel puzzle reappears.

A final set of hypotheses we collect under the umbrella of state-dependent utility. By this we mean that the representative investor in the S&P 500 index has state-dependent utility. We do

not require a representative agent for the economy as a whole. Instead, we assume the existence of the representative or marginal investor in the S&P 500 index and index options.

State-dependent utility may arise for a number of reasons. One of these is generalized utility, by which we mean that the investor exhibits either habit formation or recursive utility. For example, in the case of habit formation, and unlike with the simple power utility of Hansen and Singleton (1983), the pricing kernel m is a function of an aggregator of historical levels of consumption as well as current consumption. In this setting, an investor becomes accustomed to a level of consumption, which is determined by the aggregator, and strongly dislikes decreases in consumption below that level. Epstein and Zin (1989 and 1991) explore the recursive utility model. Here, the uncoupling of the risk-aversion coefficient and the intertemporal substitution rate is not too helpful for our purposes, since it leaves the power utility formulation across wealth in place.[13] However, Figure 3 is in violation of such power utility.

State-dependent utility may also arise when index volatility, or more generally the distribution of other security returns, is stochastic. In this case, measures of volatility, interest rates, or the prices of other assets enter the index-investor's utility as additional state variables.

We do not detail here every alternative state variable we must consider. It is possible, however, to characterize generally the nature of the relation between the index and the additional state variable that must exist to explain the puzzle. The kernel m is the ratio of expected future marginal utility divided by the current marginal utility. If utility is dependent on a second state variable, and that variable is responsible for the puzzle, then the relation between the additional state variable and the index must be non-monotonic.

To demonstrate this, consider the optimal consumption (C*) induced by a level of the S&P 500 index w and a state variable v. We then write

$$m(w) = E_v \left\{ U^{'}(C^*(w,v)) \middle| w \right\} / U^{'}(C_0^*(w_0, v_0)) \tag{4}$$

The numerator is the expected marginal utility at the time of the option expiration, where the expectation is conditioned on the index level. For equation (4) to behave as in Figure 3, the

---

[13] This feature unfortunately only became apparent after we implemented recursive utility into our model below.

distribution of v must change with w so that the numerator rises monotonically in the middle range, i.e. where w is close to w_o. However, the relation between w and the v distribution must be decidedly different in the tails. In the tails we expect that the distribution of v changes very little with w. Because marginal utility declines with wealth ceteris paribus, the numerator declines with w either far above or below w_0. We now examine several alternative representations of state-dependent utility with the goal of explaining Figure 3.

We depict a hypothetical situation of such a state variable v in Figure 5. Initially, for low wealth levels, the economy is in state 1. As wealth approaches the level of 1, state 2 becomes more likely. However, as wealth increases even further, state 1 is again becoming more likely. Thus, the pricing kernel $m*$ is a weighted average of the two pricing kernels associated with the two states of v. One can think about the following potential explanation: For very low and very high levels of wealth, utility is based on the low marginal utility function (state 1) associated with a high volatility regime. For moderate levels of wealth (close to the initial wealth level of 1), utility is based on the high marginal utility function (state 2) of a low volatility regime. Around the center, the weights shift and the resulting marginal utility function from the above Figure 3 is obtained.

Figure 5. Hypothetical Pricing Kernels depending on a Second State Variable v. We graph the simplest setting where the second state variable v can take either of two values. As wealth increases, the likelihood of being in state 2 increases. Taking the expectation over v yields the desired empirical pricing kernel m*.
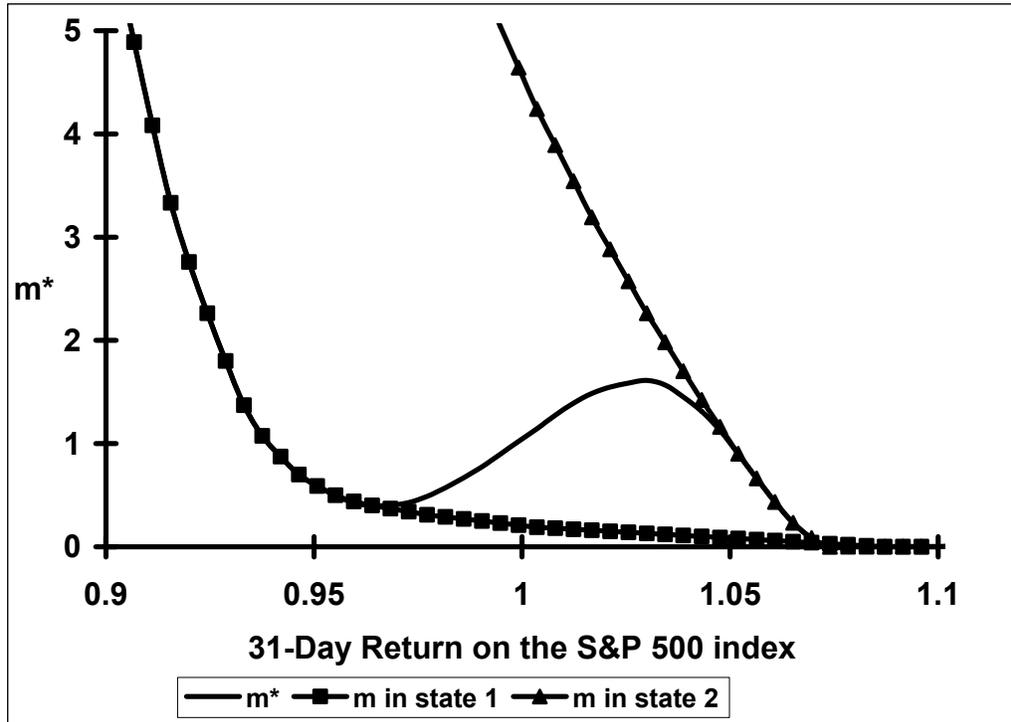
Section 3 - A Model

For our model, we stay with our above motivation and propose that volatility is a function of a second momentum state variable. This choice is based on the observation that a candidate state variable would have to exhibit an effect on the pricing kernel which is symmetric in wealth. Namely, the pricing kernel ought to be raised for small changes in wealth and lowered for large changes in wealth in either direction. Volatility naturally has this feature in that it tends to be low if wealth changes little and increases if wealth changes dramatically in either direction. Moreover, we know that in a representative agent economy with power utility and a small risk aversion coefficient, low volatility will cause the pricing kernel to be high and high volatility will lower the pricing kernel, just as desired. Finally, we are reluctant to model volatility as a function

of wealth straightaway since wealth is a non-stationary variable. Thus, we propose to model volatility as a function of a (stationary) momentum variable X_t. X_t follows a simple mean reverting process; it is an exponentially-weighted sum of past increments in wealth.

The agent's problem can then be modeled along the lines of Merton (1969) as follows. Let the representative agent choose the rate of consumption $C(W_t, X_t, t)$, and the proportion of wealth in the risky asset $\alpha(W_t, X_t, t)$, and the proportion in the riskless asset $(1-\alpha)$ to maximize expected utility of lifetime consumption:

$$J(W_t, X_t, t) \equiv \underset{\{C_s, \alpha_s\}}{Maximize} E_t \left[\int_t^\infty \exp(-\rho s) C_s^{1-B} / (1-B) \, ds \right] \tag{5}$$

where

$$dW_t = \left(W_t \alpha_t (\mu - r_t) + W_t r_t - C_t \right) dt + \alpha_t W_t \sigma(X_t) dz_t$$
$$dX_t = -\theta X_t dt + \sigma(X_t) dz_t \tag{6}$$

Here, μ and θ are constants, $\sigma(X_t)$ is a function of the second state variable X_t, and a single Brownian motion appears in both diffusion equations. Changes in the state variable are perfectly correlated with changes in wealth. Using Ito's lemma,

$$dJ = J_W dW + J_X dX + J_t dt + J_{WW} dW^2 / 2 + J_{WX} dW dX + J_{XX} dX^2 / 2 \tag{7}$$

Optimal consumption and investment choices satisfy the HJB equation:

$$0 = \underset{C_t, \alpha_t}{Maximize} \ e^{-\rho t} C_t^{1-B} / (1-B) + E_t [dJ] / dt$$
$$= \underset{C_t, \alpha_t}{Maximize} \ e^{-\rho t} C_t^{1-B} / (1-B) + J_W \left(W_t \alpha_t (\mu - r_t) + W_t r_t - C_t \right) + J_X (-\theta X_t) \tag{8}$$
$$+ J_t + J_{WW} \alpha_t^2 W_t^2 \sigma^2(X_t) / 2 + J_{WX} \alpha_t W_t \sigma^2(X_t) + J_{XX} \sigma^2(X_t) / 2$$

From the first order conditions:

$$C_t = e^{-\rho t/B} J_W^{-1/B}$$

$$\alpha_t = \frac{-J_W}{J_{WW}W_t}\frac{\mu - r_t}{\sigma^2(X_t)} - \frac{J_{WX}}{J_{WW}W_t} \tag{9}$$

In order to find the solution of the model in equilibrium, we posit the following functional form for the utility of wealth of a representative agent,

$$J(W_t, X_t, t) = e^{-\rho t} h(X_t) W_t^{1-B} /(1-B) \tag{1}$$

Substituting partial derivatives (details are in the appendix) into equation (9) we obtain

$$\frac{\partial}{\partial C} = 0 = C^{-B} - h(X_t)W^{-B} \Leftrightarrow C_t = h(X_t)^{-1/B} W_t \tag{2}$$

$$\alpha_t = \frac{\mu - r_t}{B\sigma^2(X_t)} + \frac{h'}{Bh}$$

When we assume that the riskless asset is in zero net supply, so that $\alpha_t = 1$, the latter equation identifies the riskless rate:

$$r_t = \mu - B\sigma^2(X_t) + \frac{h'(X_t)}{h(X_t)}\sigma^2(X_t) \tag{3}$$

Now, we use equation (11) to identify consumption and investment, equation (12) to identify the riskless rate, and equation (10) to identify the partial derivatives of J. Equation. (8) is written:

$$0 = (1-B)\mu h - \rho h + Bh^{1-1/B} - \theta X_t h' + \sigma_t^2(X_t)\left[\frac{1}{2}h'' + (1-B)h' - (1-B)Bh/2\right] \tag{4}$$

This identifies h without knowledge of the riskless rate. The h-function is a second order ODE with a hyperbolic term that makes it very difficult to solve. However, from the functional form of the pricing kernel we can learn much about the functional form of h. The pricing kernel is given by the ratios of marginal utilities:

$$m_{t,T} = \frac{J_{W_T}}{J_{W_t}} = e^{-\rho(T-t)}\left[\frac{h(X_T)}{h(X_t)}\right]\left(\frac{W_T}{W_t}\right)^{-B} = e^{-\rho(T-t)}\left[\frac{h(X_T)}{h(X_t)}\right]r_{m,t,T}^{-B} \qquad (5)$$

The adjustment depends on the ratio of $h(X_T) / h(X_t)$ which will give the desired shape of the pricing kernel if h is shaped like a normal density (Bell-curve) but raised by a constant above the zero line. We thus specify the following parametric form for h:

$$h(X) = d + e^{a+bX+cX^2} \qquad (6)$$

We are now in the position to look at results by choosing some reasonable set of parameters. We initialize our economy with unit wealth $(W_0 = 1)$ and zero momentum $(X_0 = 0)$. The risk aversion coefficient of the representative agent is set to $B = 0.5$ and the mean reversion of the momentum process is fixed at $\theta = 0.1$. We specify the adjustment function as:

$$h(X) = 8 + e^{-500X^2} \qquad (7)$$

and the volatility function as:

$$\sigma(X) = 0.06 + 0.2X^2 \qquad (8)$$

We are then able to derive the drift μ and the risk free rate r from equations (13) and (12), respectively. This allows us to simulate the wealth and the momentum processes based on a daily discretization over 31 days. We use 10,000 simulation runs and smooth the resulting graph through a kernel regression. We document in Figure 6 the simulated pricing kernel which indeed is locally increasing in wealth around a wealth level of 1. It is by design still decreasing in consumption but the second state variable – in our case momentum – introduces a wedge

between consumption and wealth. In comparison with Figure 3 we notice that the model generates a much less pronounced hump in the pricing kernel. This is caused by the economics of the underlying model in which the wedge between consumption and wealth cannot be made arbitrarily large. As a result, the actual and the risk-neutral probability distribution still look more alike than their empirical estimates in Figure 2. We find that consumption is on the low side at a annualized rate of 1.23-1.56% of wealth.

Figure 6. Simulated Pricing Kernel based on a Representative Agent Model with a Second State Variable. 10,000 Simulation runs and smoothed through a kernel regression.
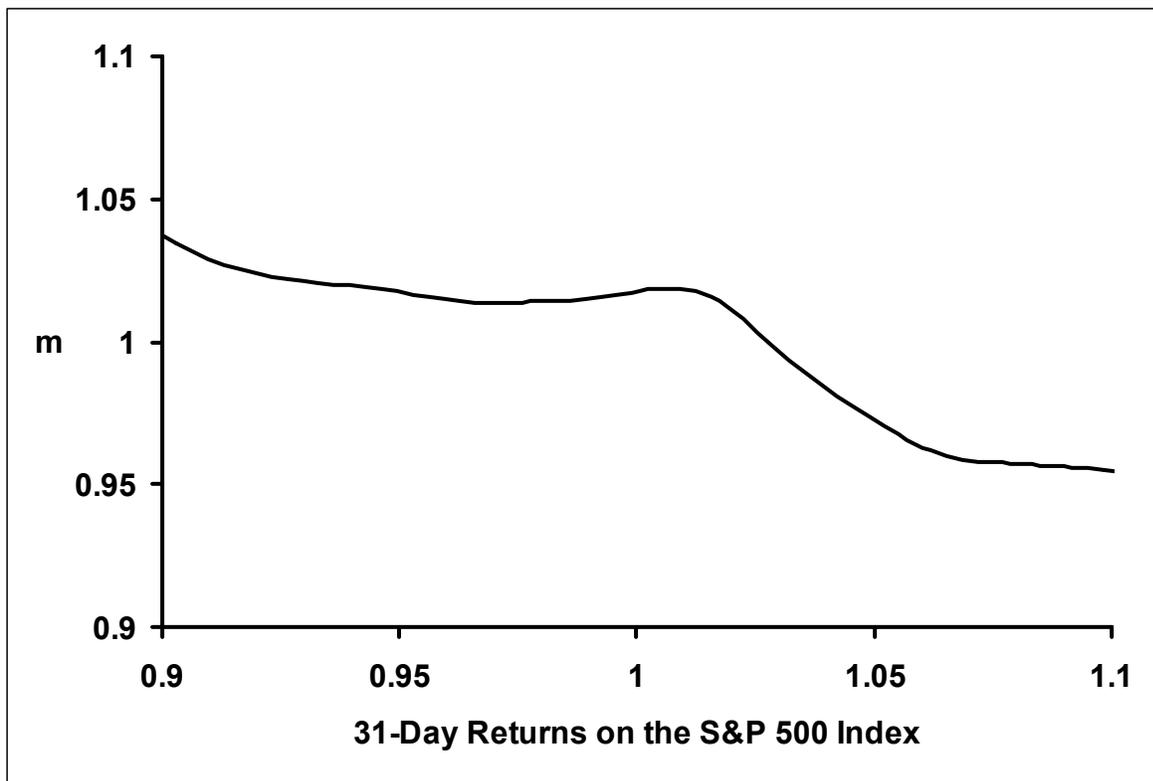

The realized volatility of the wealth process mimics the u-shaped specification of volatility as a function of momentum (X). We see in Figure 7 that the realized volatility is also u-shaped when we plot it across 31-day returns and smooth the graph through a kernel regression again. The far edges of the volatility graph ramp up to rather high values.

Figure 7. Annualized Volatility of Simulated Wealth Process based on a Representative Agent Model with a Second State Variable. 10,000 Simulation runs and smoothed through a kernel regression.
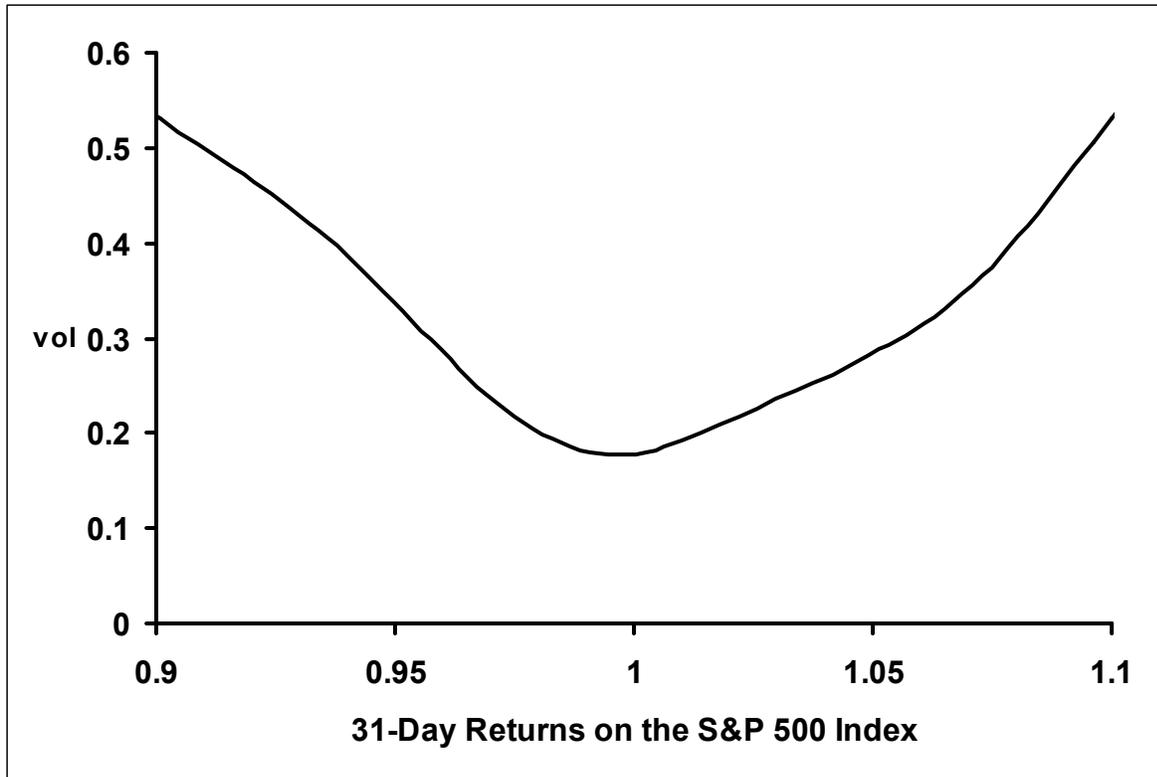
We also find that the model implied function of the drift and the riskless rate change very rapidly as a function of momentum. This is due to the high curvature in the adjustment function around a wealth level of 1. However, when we consider the drift and the riskless rate across wealth, the resulting functions are much better behaved as we show in Figure 8. The annualized unconditional drift is 8.68% and the annualized unconditional riskless rate is 7.28%. The riskfree rate is thus a little high and the equity risk premium is rather small compared to the historical risk premium. However, that we are faced with the equity premium puzzle is not surprising for this type of model.

Figure 8. Annualized Drift and Riskless Rate of Simulated Wealth Process based on a Representative Agent Model with a Second State Variable. 10,000 Simulation runs and smoothed through a kernel regression.
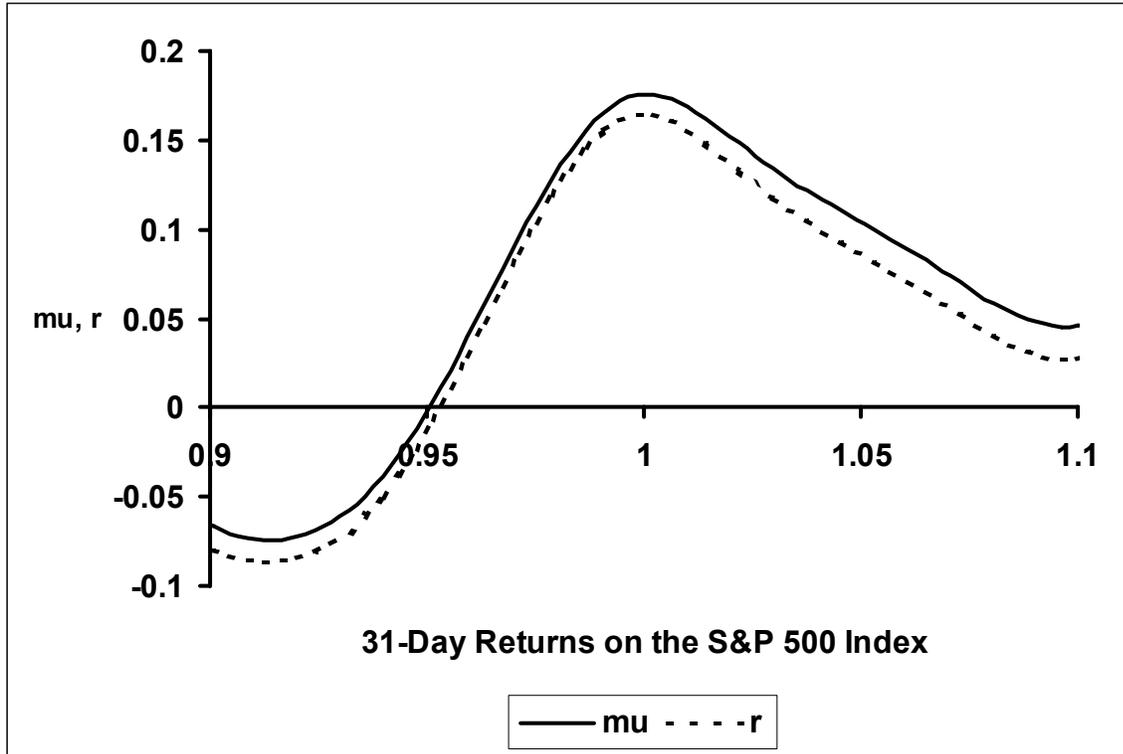
Section 4 – Summary

The goal of our work is to explain the pricing kernel puzzle pictured in Figure 3. This figure demonstrates a surprising pattern of the marginal utility of wealth for investors in the S&P 500 index. It appears that these investors are risk seeking, at least with respect to small risks in the index. This is inconsistent with the basic assumptions of asset-pricing theory, when it is recognized that the S&P 500 index represents a significant proportion of public equity capital in the U.S.

Our theoretical model provides for a representative agent model that features volatility as a function of an additional state variable which measures momentum. Through an adjustment function to the pricing kernel, the model is capable of generating the sort of relationship we see in Figure 3. Future research will examine index returns from markets outside the U.S along the

26

lines of Jackwerth (2004). His limited sample seems to suggest that the puzzle is not unique to the S&P 500 index options market but exists internationally. Further empirical work could then use both domestic and foreign market data to test our theoretical explanation of the puzzle.

Appendix:

Formulae

Write the utility of wealth of a representative agent as:

$$J(W_t, X_t, t) = e^{-\rho t} h(X_t) W_t^{1-B} / (1-B),$$ (A1)

so the derivatives of utility of wealth are

$$J_W = e^{-\rho t} h(X_t) W_t^{-B}$$

$$J_t = -\rho e^{-\rho t} h(X_t) W_t^{1-B} / (1-B)$$

$$J_x = e^{-\rho t} h'(X_t) W_t^{1-B} / (1-B)$$ (A2)

$$J_{xx} = e^{-\rho t} h''(X_t) W_t^{1-B} / (1-B)$$

$$J_{ww} = e^{-\rho t} (-B) h(X_t) W_t^{-1-B}$$

$$J_{wx} = e^{-\rho t} h'(X_t) W_t^{-B}$$

The derivation of Eq. 13 out of Eq. 8 uses the optimal $\alpha_t = 1$ and proceeds through:

$$0 = e^{-\rho t} \left(h(X_t)^{-1/B} W_t \right)^{1-B} / (1-B)$$

$$+ (e^{-\rho t} h(X_t) W_t^{-B}) \left(W_t (\mu - r) + W_t r_t - \left(h(X_t)^{-1/B} W_t \right) \right)$$

$$+ (e^{-\rho t} h'(X_t) W_t^{1-B} / (1-B))(-\theta X_t) - \rho e^{-\rho t} h(X_t) W_t^{1-B} / (1-B)$$ (A3)

$$+ (e^{-\rho t} (-B) h(X_t) W_t^{-1-B}) W_t^2 \sigma^2 (X_t) / 2$$

$$+ (e^{-\rho t} h'(X_t) W_t^{-B}) W_t \sigma^2 (X_t) + 0.5 e^{-\rho t} \sigma^2 (X_t) h''(X_t) W_t^{1-B} / (1-B)$$

The empirical tests are based on a database containing all minute-by-minute European option quotes and trades on the S&P500 index from April 2, 1986 to December 29, 1995. We use only option quotes since we cannot know for actual trades where they occurred relative to the bid/ask spread and our results might be affected. The database also contains all futures trades and quotes on the S&P 500. Our goal is to obtain a panel of daily return observations on the index, the risk-free rate, and on several options with different strike price/index level ratios (moneyness) and constant maturity.

Index Level. Traders typically use the index futures market rather than the cash market to hedge their option positions. The reason is that the cash market prices lag futures prices by a few minutes due to lags in reporting transactions of the constituent stocks in the index. We check this claim by regressing the index on each of the first twenty minute lags of the futures price. The single regression with the highest adjusted R^2 was assumed to indicate the lag for a given day. The median lag of the index over the 1542 days from 1986 to 1992 was seven minutes. Because the index is stale, we compute a future-based index for each minute from the future market

$$S_t = \left(\frac{r}{d}\right)^{-\Delta} F_{t+\Delta} \tag{A4}$$

For each day, we use the median interest rate implied by all futures quotes and trades and the index level at that time. We approximate the dividend yield by assuming that the dividend amount and timing expected by the market were identical to the dividends actually paid on the S&P 500 index. However, some limited tests indicate that the choice of the index does not seem to affect the results of this paper.

Interest Rates. We compute implied interest rates embedded in the European put-call parity relation. Armed with option quotes, we calculate separate lending and borrowing interest returns from put-call parity where we used the above future-based index. We assign, for each expiration date, a single lending and borrowing rate to each day, which is the median of all daily observations across all striking prices. We then use the average of those two interest rates as our daily spot rate for the particular time-to-expiration. Finally, we obtain the interpolated interest

rates from the implied forward curve. If there is data missing, we assume that the spot rate curve can be extrapolated horizontally for the shorter and longer times-to-expiration. Again, some limited tests indicate that the results are not affected by the exact choice of the interest rate.

Options with adjusted Moneyness and constant Maturity. It is important to use options with adjusted moneyness and constant maturity since our test statistics involve the conditional covariance matrix of option pricing errors. If the maturity of the options were not constant over time, then the conditional covariance matrix of the pricing errors would be time varying, too. This would require additional exogenous assumptions on the structure of the covariance matrix and the estimation of several additional parameters, which could lead to additional estimation error in our test statistics.

In our data set, all puts are translated into calls using European put-call parity. Then, we compute the implied volatilities where we use the Black-Scholes formula as a translation device only. We then adjust throughout each day for the movement of the stock price by assuming that the implied volatilities are independent of the underlying stock price. Then, we pick the stock price closest to 12 pm as our daily stock price and value all options from throughout the day as if they were call options with the implied volatilities measured above and struck at the moneyness level measured above. We do not eliminate any daily observations due to their level of moneyness.

Arbitrage Violations. In the process of setting up the database, we check for a number of errors, which might have been contained in the original minute-by-minute transaction level data. We eliminate a few obvious data-entry errors as well as a few quotes with excessive spreads -- more than 200 cents for options and 20 cents for futures. General arbitrage violations are eliminated from the data set. We also check for violations of vertical and butterfly spreads. Within each minute we keep the largest set of option quotes which does not violate:

$$Sd^{-t} \geq C_i \geq \max[0, Sd^{-t} - K_i r^{-t}] \qquad (A5)$$

American early exercise is not an issue as the S&P500 options are European in nature, and the discreteness of quotes and trades only introduces a stronger upward bias in the midpoint implied volatilities for deep-out-of-the-money puts (moneyness less than 0.6) which we do not use in our empirical work.

References

Agarwal, Vikas, and Narayan Y Naik (2000). "Performance Evaluation of Hedge Funds with Option-based and Buy-and-Hold Strategies", working paper, London Business School.

Bakshi, G., C. Cao, and Z. Chen (1997). "Empirical Performance of Alternative Option Pricing Models." *Journal of Finance* 52, No. 5, 2003-2049.

Bartunek, Kenneth, and Mustafa Chowdhury (1997). "Implied Risk Aversion Parameter from Option Prices." *The Financial Review* 32, No. 1, 107-124.

Bates, David (2000). "Post-'87 Crash Fears in S&P 500 Futures Options," *Journal of Econometrics* 94, No. 1-2, 181-238.

Benninga, Simon, and Jarom Mayshar (1997). "Heterogeneity and Option Pricing", working paper, University of Pennsylvania.

Black, Fisher (1976). "Studies of Stock Price Volatility Changes", Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economical Statistics Section, 177-181.

Bliss, Robert R, and Nikolaos Panigirtzoglou (2004). "Option-Implied Risk Aversion Estimates." *Journal of Finance*, Feb, forthcoming

Bondarenko, Oleg (2003). "Why are Put Options so Expensive?" Working Paper, University of Illinois at Chicago November 2003.

Branger, Nicole, and Christian Schlag (2002). "Why is the Index Smile so Steep?" Working Paper, University of Frankfurt.

Breeden, D. and R. Litzenberger (1978). "Prices of State-Contingent Claims Implicit in Options Prices." *Journal of Business* 51, 621-651.

Brown, David P., and Michael R. Gibbons (1985). "A Simple Econometric Approach for Utility-Based Asset Pricing Models." *Journal of Finance* 40, 359-381.

Buraschi, Andrea, and Jens C. Jackwerth (2001). "The Price of a Smile: Hedging and Spanning in Option Markets", *Review of Financial Studies* 14, No. 2, 495-527.

Campa, Jose M., and P.H. Kevin Chang (1998). "ERM Realignment Risk and its Economic Determinants as Reflected in Cross-Rate Options." *Economic Journal*, 1046-1066.

Campbell, John Y, and John H Cochrane (2000). "Explaining the Poor Performance of Consumption-Based Asset Pricing Models," *Journal of Finance* 55, No. 6, 2863-2878.

Campbell, John Y., Martin Lettau, Burton G. Malkiel, and Yexiao Xu (2000). "Have Individual Stocks Become More Volatile? An Empirical Exploration of Idiosyncratic Risk", *Journal of Finance* 56, No. 1, 1-43.

Christie, A. (1982). "The Stochastic Behavior of Common Stock Variance – Value, Leverage and Interest Rate Effects", *Journal of Financial Economics* 10, 407-432.

Coval, Joshua D, and Tyler Shumway 2001 Expected Option Returns Journal of Finance 56, No. 3, June 2001, 983-1009

Cox, John C., Jonathan E. Ingersoll, and Stephen A. Ross (1985). "An Intertemporal General Equilibrium Model of Asset Prices." *Econometrica*, 363-384.

Debreu, Gerard. (1993). "Theory of Value: An Axiomatic Analysis of Economic Equilibrium." Yale University Press, reprinted, 114 pp.

Dennis, Patrick, and Stewart Mayhew (2000). "Implied Volatility Smiles: Evidence From Options on Individual Equities." Working Paper, University of Virginia.

Driessen, Joost, and Pascal Maenhout (2003). "A Portfolio Perspective on Option Pricing Anomalies." Working Paper, University of Amsterdam.

Duan, Jin-Chuan (1995). "The GARCH Option Pricing Model". *Mathematical Finance* 5, No. 1, 13-32.

Epstein, Larry G and Stanley E Zin (1989). "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework". *Econometrica* 57, No. 4, 937-969.

Epstein, Larry G and Stanley E Zin (1991). "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis". *Journal of Political Economy* 99, No 2, 263-286.

Fama, Eugene F. and Kenneth R. French (1988). "Permanent and Temporary Components of Stock Prices." *Journal of Political Economy*, 246-273.

Franke, Günter, Dick Stapleton, and Marti Subrahmanyam (1998). "Who Buys and Who Sells Options: The Role of Options in an Economy with Background Risk", *Journal of Economic Theory* 82, No. 1, 89-109.

Franke, Günter, Dick Stapleton, and Marti Subrahmanyam (1999). "When are Options Overpriced? The Black-Scholes Model and Alternative Characterisations of the Pricing Kernel", working paper, Konstanz University.

Gennotte, Gerard and Hayne E. Leland (1990). "Market Liquidity, Hedging, and Crashes", *American Economic Review* 80, Ro. 5, 999-1021.

Grossman, Sanford J. (1988). "An Analysis of the Implications for Stock and Futures Price Volatility of Program Trading and Dynamic Hedging Strategies", *Journal of Business* 61, 275-298.

Grossman, Sanford J. and Zhongquan Zhou (1996). " Equilibrium Analysis of Portfolio Insurance." *Journal of Finance*, 1379-1403.

Hansen, Lars P. and Kenneth J. Singleton (1983). "Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns." *Journal of Political Economy*, 249-265.

He, Hua, and David M. Modest (1995). "Market Frictions and Consumption-Based Asset Pricing", *Journal of Political Economy* 103, No. 1, 94-117.

Hentschel, Ludger (2002). "Errors in Implied Volatility Estimation." *Journal of Financial and Quantitative Analysis*, forthcoming.

Hong, Harrison, and Jeremy C. Stein (1999). "Differences of Opinion, Rationale Arbitrage and Market Crashes", working paper, Stanford University.

Jacklin, Charles J., Allan W. Kleidon, and Paul Pfleiderer (1992). "Underestimation of Portfolio Insurance and the Crash of October 1987", *Review of Financial Studies* 5, No. 1, 35-63.

Jackwerth, Jens C. (2000). "Recovering Risk Aversion from Option Prices and Realized Returns", *Review of Financial Studies* 13, No. 2, 433-451.

Jackwerth, Jens C. (2004). "Option Implied Risk-Neutral Distributions and Risk Aversion." Monograph for AIMR, forthcoming.

Jackwerth, Jens C. (1999). "Option Implied Risk-Neutral Distributions and Implied Binomial Trees: A Literature Review," *Journal of Derivatives* 7, No. 2, 66-82.

Jackwerth, Jens C., and Mark Rubinstein (1996). "Recovering Probability Distributions from Option Prices," *Journal of Finance* 51, No. 5, 1611-1631.

Kelly, M. (1994). "Correlation: Stock Answer", *RISK* 7, No. 8, 40-43.

Kleidon, Allan W. (1994). "Stock Market Crashes", working paper, Stanford University.

Lucas, Robert E., Jr., "Asset Prices in an Exchange Economy." *Econometrica* (1978), 1429-1445

Merton, Robert (1969). "Lifetime Portfolio Selection Under Uncertainty: The Continuous-Time Case," *Review of Economics and Statistics* 51, 247-257.

Nandi, Saikat (1999). "Asymmetric Information about Volatility: How Does it Affect Implied Volatility, Option Prices and Market Liquidity?" *Review of Derivatives Research* 3, No. 3, 215-236.

Pan, Jun (2002). "The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study," *Journal of Financial Economics* 63, 3-50.

Poterba, James M. and Lawrence H. Summers (1988). "Mean Reversion in Stock Prices," *Journal of Financial Economics*, 27-59.

Romer, David (1993). "Rational Asset-Price Movements without News", *American Economic Review* 83, 1112-1130.

Rosenberg, Joshua V., and Robert F. Engle, (2002). "Empirical Pricing Kernels." *Journal of Financial Economics* 64, No. 3, forthcoming.

Rubinstein, Mark (1994). "Implied Binomial Trees," *Journal of Finance* 49, No. 3, 771-818.

Rubinstein, Mark (1976). "The Valuation of Uncertain Streams and the Pricing of Options," *Bell Journal of Economics* 407-425.

Schwert, G. William (1989). "Why Does the Stock Market Volatility Change over Time?" *Journal of Finance* 44, 1115-1153.

Toft, Klaus and B. Prucyk (1997). "Options on Leveraged Equity: Theory and Empirical Tests," *Journal of Finance* 52, No. 3, 1151-1180.

Ziegler, Alexandre (2003). "Why does Risk Aversion Smile?" Working Paper, University of Lausanne.