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 18, 2002

C:\Research\Paper18\paper26.doc

Abstract

One of the central questions in financial economics is the determination of asset prices, such as the value of a stock. Over the past three decades, research on this topic has converged on a concept called the "state-price density". However, a puzzle has arisen. On the one hand, Cox, Ingersoll, and Ross (1985) and others argue that the ratio of the state-price density to the statistical probability density, which is commonly known as the pricing kernel, should decrease monotonically as the aggregate wealth of an economy rises. On the other hand, recent empirical work on options on the S&P 500 index suggests that, for a sizable range of index levels, the pricing kernel is increasing instead of decreasing.

We investigate theoretical explanations to this puzzle. Our existing work has ruled out some alternative hypotheses, such as data imperfections and methodological problems. We provide a 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. We estimate the model through GMM.

* 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 and Mark Rubinstein 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

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. Furthermore, as Bernstein (1993) has argued, these inventions and the work that followed them were critical in the development of modern financial securities, methods of risk management, and the structure of today's financial markets.

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).

Because of its central role in pricing assets, the state-price density now has many uses. For example, central banks use estimates of SPDs to measure the credibility of exchange rate commitments.[1] Another application 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 agent's optimality conditions 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.[2] 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 SPD (discussed by Cox, Ingersoll, and Ross (1985)) onto the space of payoffs defined by the index. In

[1] See Campa and Chang (1998) and the survey by Jackwerth (1999).

[2] For example, some central banks draw their estimates of the state-price density from foreign exchange rates and the prices of foreign exchange options.

simple terms, this means 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. More generally, a SPD can be used to price all derivative securities based on an underlying asset, when options on the underlying are used to estimate the SPD.

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. 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 sugggest. 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. Section 2 links the discussion of the pricing kernel to the option smile. In section 3 we discuss our empirical investigation of several hypotheses regarding our estimates of marginal utility and their positive relation to the index. 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,

Eq. 1
$$C = \sum_s p_s F_s \, ,$$

where F_s is the option payoff and p_s is the current Arrow-Debreu state price of a dollar paid in the future state s.[3,4] The sum of the p_s across states is equal to one, so the set of p_s represent a distribution analogous to, but distinct from a market participant's subjective distribution of probabilities of the states.[5] The set of p_s is a state-price density or SPD.[6]

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 p_s satisfies

Eq. 2 $$p_s / q_s = m_s$$

where q_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 probability is high. Conversely, if one observes a state price far below a probability belief, 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

Eq. 3 $$C = \sum_s m_s q_s F_s = \sum_s \left(U'\left(C_s\right) / U'\left(C_o\right) \right) q_s F_s \ ,$$

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 implication of power utility is that aggregate consumption is proportional to aggregate wealth.[7] 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 p and q using data on the S&P 500 index return and option prices from April 2, 1986 to the end of 1995.[8] 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 used, and this method does not require a representative agent

---

[3] See Debreu (1993) for development of a state price.

[4] 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, and a strike price defined in the option contract. We write $F_s = \text{Max}[0, P_s - X]$.

[5] 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 in (1). In general, the price of a security is the sum of expected future payoffs (using the SPD) discounted at the risk-free rate.

[6] The word "density" is usually reserved for a continuous function. For simplicity of discussion, we refer to the probability distribution in (1) as a state-price density.

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

[8] Rosenberg and Engle (1999) confirm this shape by using a related method.

to be specified. An estimate of p is obtained by the method applied to the option data. Multiple options with various strike prices, but all with 31-day lives are used. The subjective probability distribution q is estimated as the historical distribution of index returns over 31-day intervals. This is an appropriate method for estimating q provided that investors have rational expectations. Typical estimates of the distributions, say p^* and q^*, are graphed in Figure 1.

Figure 1. 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.

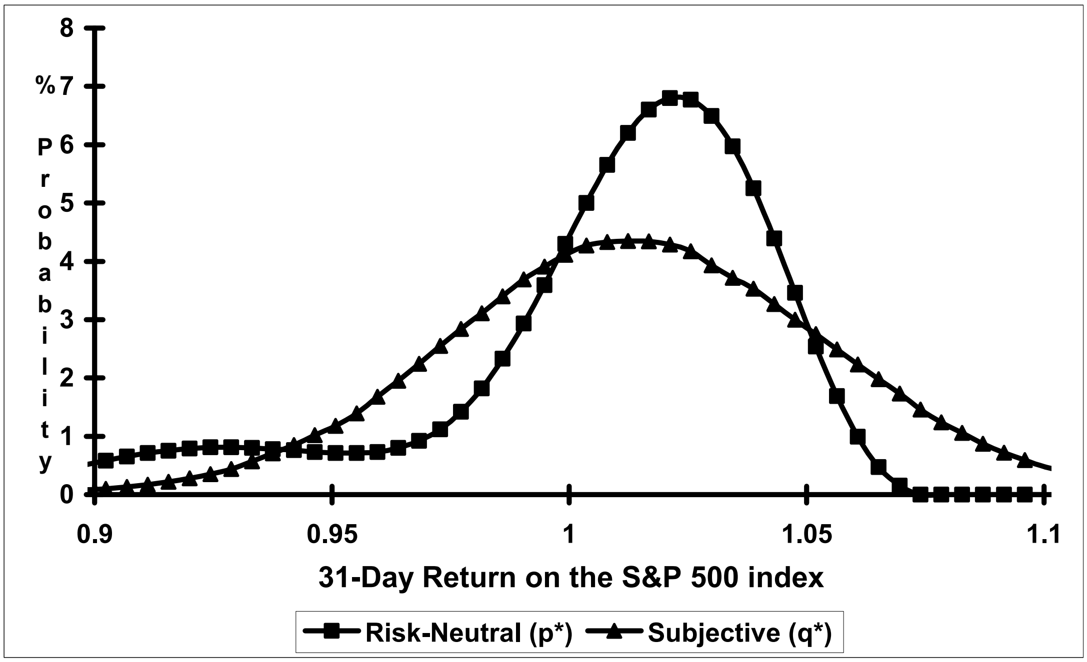

A typical estimate of m, say m^*, appears in Figure 2 as a function of the return on the S&P 500 index. The value of one in the center of the horizontal axis represents an ending level of the index (i.e., at the time of option expiration) that is equal to the current level (at the beginning of the 31-day interval). Globally, m^* is a decreasing function of the ending 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 occurs because the proportional difference between the estimates, say p^* and q^*, is increasing over this range, as equation (2) implies.

Figure 2. 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.

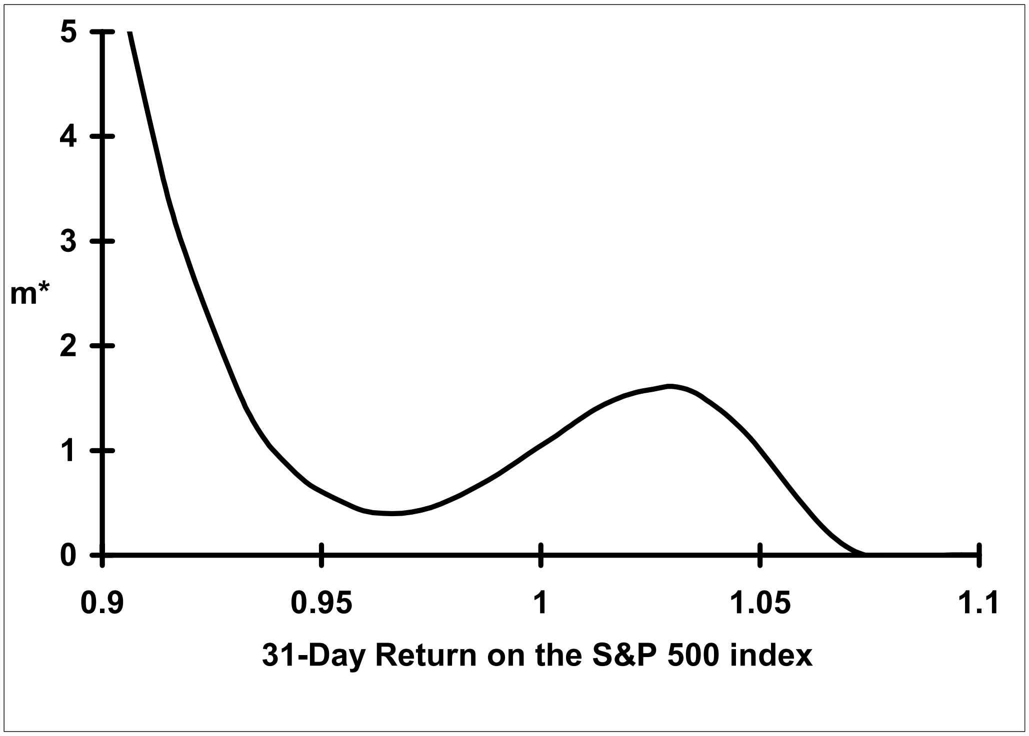

Traditional asset pricing theory, e.g. Rubinstein (1976) and Lucas (1978), assumes that a representative investor exists. It also is common in tests of asset pricing theories to assume that a market index such as the S&P 500 index represents the aggregate wealth held by this investor. If we make these assumptions, then the empirical m^* of Figure 2 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 2 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 p Jackwerth and Rubinstein (1996) only require that the prices do not offer any arbitrage trading opportunities, while the sample estimator of q requires rational expectations to be valid. Therefore we interpret m^* differently from that just given. 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 (proportional to) the marginal utilities of investors 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.

Section 2 - Option Smiles

Option prices exhibit so-called volatility smiles and skews. These labels refer to systematic patterns in the cross section of option prices. To explain this idea, we note that option prices are commonly transformed to and measured in units of volatility of the underlying. 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 a substitute for the option price. Once the transformation from price to volatility is made it is an easy task to compare the prices of options. For example, researchers often describe the cross section of options by a plot of implied volatilities against strike prices as in Figure 3.[9]

Figure 3. Empirical Volatility Smile. Typical post-1987 stock market crash implied volatility smile. Returns are reported as 1 plus the rate of return.

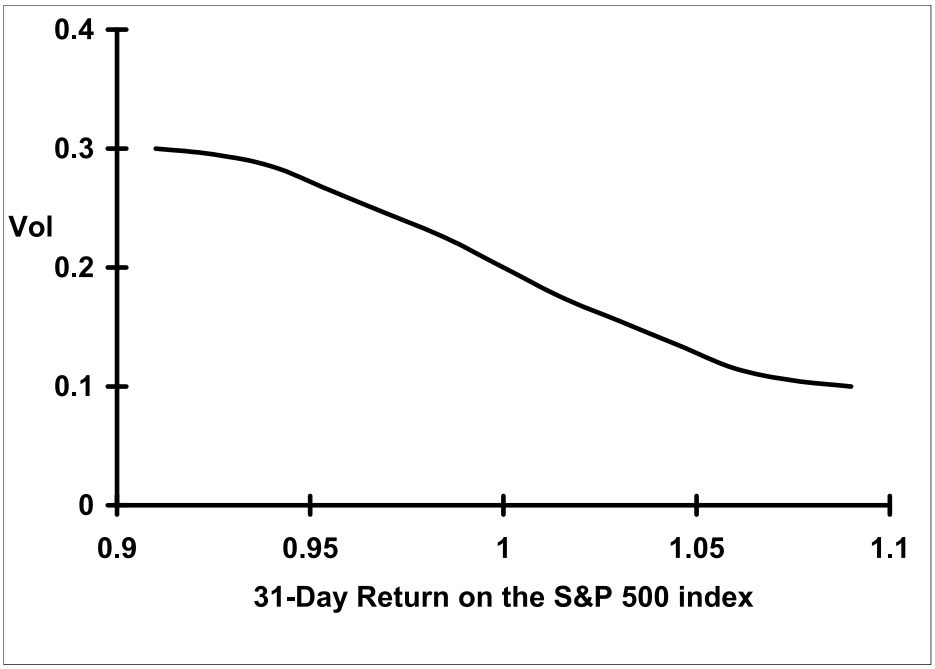

A consequence of the Black-Scholes option-pricing model is that implied volatilities of all options are equal. If this result held, the plot in Figure 3 would be flat. Instead, what we observe is a "smile", or "skew", or "smirk" depending on your preference in label. The left tail of the plot is substantially above the right. A plot such as this is a common finding, although the degree to which the left tail of the plot is elevated relative to the center is systematically related to the underlying security. For example, for options on individual equities, the plot is relatively flat compared to Figure 3. For foreign-exchange options, each tail of the plot is slightly and equally elevated. For the S&P 500 index options, the smile looks like that of Figure 3, where the left tail is substantially above the right tail.

For any option-pricing model, the plot of the kernel m (as in Figure 2) is closely tied to the degree of the smile (like that of Figure 3).[10] It is important to note that any model used to explain

[9] Pricing based on implied volatility is analogous to the use of yield to maturity in the bond market. Just as the yield to maturity is a transformation of the bond price, the implied volatility is a transformation of the option price. Implied volatilities allow traders to easily compare the prices of options with disparate terms, e.g., different strike prices.

[10] 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 (1999).

the smile in the S&P 500 index options must provide a kernel m that is consistent with Figure 2. 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 2.

Section 3 - Alternative Hypotheses

In principle, a number of alternative hypotheses might explain the increasing marginal utility shown in Figure 2. Some of these are discussed elsewhere, so we briefly review them in section 3.1. We discuss the literature in section 3.2. We propose new hypotheses to investigate and discuss these in section 3.2.

Section 3.1 – Hypotheses Discussed Elsewhere

One hypothesis that explains the puzzle shown in Figure 2 is that either one or both of our estimation procedures for p and q are faulty. A second possibility is that the option prices used to calculate p^* are noisy. Jackwerth (2000) discusses these concerns in detail and has ruled them out. Here, we briefly review his findings.

First note that Jackwerth and Rubinstein (1996) provide a number of alternative methods to obtain estimates of p. Jackwerth (2000) examines these methods using simulated data for which we know the true p, and he shows them to be quite accurate. Similarly, a variety of standard procedures, both parametric and non-parametric, are used to estimate the historical return distribution q^* and they provide equivalent results. Our work will rely on the same data, which contains records of all trades and quotes in the index options and a continuous record of the index levels, again from April 2, 1986 to 1995.[11] Although there are bid-ask spreads in options markets, midpoints of the bid-ask spread are used to calculate p^* instead of trade prices. We also use all option quotes during a day for an estimate of p^* for that day. Jackwerth (2000) demonstrates that measurement of p^* is not sensitive to the existence of spreads. Furthermore, there is no good reason to believe that our index return data is poor. In summary, neither data errors nor poor estimation techniques are the cause of the puzzle.

Another hypothesis is that the estimate m^* in Figure 2 is very noisy. Perhaps the standard error bounds on m^* are sufficiently large that the slope in the middle range (about the value of 1) is statistically insignificantly different from zero. Jackwerth (2000) rules this hypothesis out. The slope is significantly positive.

Our data includes index option prices during the months before the market crash of 1987. Estimates of pricing kernels from this period are downwards sloping, and do not exhibit the puzzling behavior of Figure 2. Estimates of m^* from data after the crash universally exhibit the puzzle. This suggests that the 1987 crash induced a radical change in investors' beliefs. Beginning in 1988 and continuing through the end of our data, p^* impounds a belief that market crashes are highly likely relative to our estimate q^*. This can be seen in the left tail of the distributions of Figure 1.

One interpretation of the pre- to postcrash comparison is the following: On the one hand, investors rationally updated their beliefs, and this caused the changed estimates of p. On the other hand, our estimate q^* suffers from a so-called "peso problem". By this we mean that even though it is an accurate estimate of the statistical distribution, q^* does not represent the true

[11] Details can be found in the appendix.

distribution faced by investors because there are too few crashes in our sample. If so, q^* does not incorporate correctly the likelihood of an infrequent event, while investors are cognizant of the correct likelihood.

Jackwerth (2000) examines the peso problem by calculating statistical distributions with simulated crashes. To the historical index-return data, daily market-crash returns (that is daily returns of –20%) are randomly mixed in. The pricing kernel is estimated using q^* derived from a mixture of historical and simulated returns. He then shows that q^* calculated in this fashion do not exhibit sufficiently fat left tails; they do not reconcile the fat left tails in the estimates of p. In other words, the resulting m^* derived in this manner is a locally increasing function of the index. We conclude that a peso problem is not the source of the puzzle.

Section 3.2 – Hypotheses in the Literature Regarding the Smile

A number of explanations have been brought forward in the literature in order to explain the volatility smile. However, we will see that the existing models in the literature are largely unsatisfactory in explaining the puzzle of Figure 2. Most of the research is one-sided, meaning that researchers either model the index return process (and thereby generate a q^* distribution, which is typically close to a lognormal) or they explain the post-crash option prices (and thereby generate a left-skewed p^* distribution), but not both at the same time. Thus, these works do not touch on the puzzle we investigate here. The few existing papers that model both q^* and p^* simultaneously circumvent the puzzle altogether. For example, Pan (1999) assumes a specific pricing kernel that is monotonically decreasing in wealth. We now proceed to survey and document the inconsistency of existing models with the form of the kernel estimate m^* shown in Figure 2.

A few of those existing models can be dismissed right away. Namely, the American nature of some options can introduce a volatility smile but this is not a concern to the present study since the S&P 500 index options are European in nature. Also, the discreteness of quotes and trades is not likely to be an issue, either. While discreteness will force the implied ask volatilities of away-from-the-money options up, it forces the bid implied volatilities for those options down. Since we use midpoint implied volatilities the effects largely cancel around at-the-money and only lead to increased implied volatilities for moneyness levels below 0.6, which are seldom traded and not used in our estimates. Thus, we have to look for other arguments for the existence of the steep index option smile.

Leverage effect: Black (1976), Christie (1982), and Schwert (1989) first advocated the leverage effect. As asset price falls, the market value debt-equity ratio rises. Then, subsequent shocks to the firm value would imply higher equity volatility. However, even for individual stock options with their comparatively flat smiles, Toft and Prucyk (1997) find that the leverage effect can only account for about half of the observed smile. Also, the actual leverage on the index should be similar to the average of leverages across all firms. But such index leverage effect would be far too small to explain the steep index smile. Surprisingly, Dennis and Mayhew (2000) find evidence of the reverse of a leverage effect. They also criticize the methodology of Toft and Prucyk (1997) in that their slope variable is contaminated by scaling by the implied at-the-money volatility, which can lead to spurious leverage effects in the regressions.

Information aggregation: Grossman (1988), Gennotte and Leland (1990), Jacklin, Kleidon, and Pfleiderer (1992), Kleidon (1994), and Romer (1993) develop models in which an investor learns about an important variable through trading. Once the learning takes place, prices adjust sharply, even if no news event took place. Unfortunately, these models are inherently symmetric in that the misperception of the investor could be an over- or under-estimation, i.e. the price adjustment could be either a crash or a "melt-up". However, large crashes are more likely empirically than large up moves. Hong and Stein (1999) develop a related model in which the pessimistic investors are short sale constrained. Their model is asymmetric in that crashes are more likely than up moves. Nandi (1999) proposes a model where a risk-neutral trader has private information about future volatility. However, the symmetric smile generated is very slight and not nearly as pronounced as the empirical smile in the S&P 500 index options market.

Correlation: Kelly (1994) argues that in down-markets correlations across companies increase. This effect exacerbates the impact of bad news on the market and can lead to increased volatility for the index. The explanation for this behavior is that equities become more highly correlated in down markets since disaster affects everybody. However, good news is not as much correlated across companies and volatility rises therefore only during market downturn but not during up-moves. Kelly (1994) and Campbell et al (2001) find that individual stocks are more risky today than they were in 1963 or even in 1986. Conversely, the stock market as a whole did not show such trend in volatility. The explanation for this behavior lies in the correlations across stocks that decreased over this time. However, the index volatility smile has not flattened out over the 10 years after the crash. Moreover, the index smile changed dramatically on the day of the crash, and the correlation structure of returns developed much more gently over time.

Wealth effect and utility functions: As the market falls, the wealth of investors falls, and investors become more risk-averse. Then, the same news leads to a greater reaction and to more trading in order to rebalance the optimal portfolio for the investor. This in turn causes volatility to rise. Again, only modest smiles can be achieved as long as the associated pricing kernel m is monotonically decreasing across wealth. Also, the pricing kernel cannot be chosen independently of the stochastic process of the underlying, and we need to address the relationship in the next subsection. We will now turn to some examples of more flexible pricing kernel, which can generate (small) smiles.

Franke, Stapleton, and Subrahmanyam (1999) suggest a pricing kernel with declining elasticity. Such pricing kernel is consistent with the concept of the investors being faced with undiversifiable background risk, Franke, Stapleton, and Subrahmanyam (1998). However, their model does not imply an increase in marginal utility over the near-the-money range, which we find empirically in Figure 2.

Benninga and Mayshar (1997) obtain moderate smiles by using a more flexible pricing kernel. In particular, even in their richer setup of heterogeneous investors who can choose between current and future consumption, the marginal utility patterns of Figure 2 cannot be achieved unless a) the subjective probability assessment of future wealth is not nearly lognormal but severely left-skewed and leptokurtic or b) individual marginal utilities are increasing as consumption increases from slightly below current consumption to slightly above current consumption.

Campbell and Cochrane (2000) explore utility functions with habit persistence. They reset the starting point of the wealth dimension and over time, the investor exhibits risk-aversion increase for certain wealth (consumption) levels as the habit level increases. However, the reference point is just one point indeed, and it is known at all times. Given the reference point, the utility function is still a power utility, and the monotonicity of marginal utility is preserved. Looking forward to a solution of the puzzle, we would need the reference point to vary in conjunction with wealth (consumption), thus acting as an additional state variable. We will further explore this avenue below.

Volatility feedback: The volatility feedback effect argues that some large news event drives up the volatility of the underlying and the risk premium rises. If the news event was positive (negative), then the resulting up (down) move in the index will be dampened (amplified). Pindyck (1984), French, Schwert, and Stambaugh (1987), and Campbell and Hentschel (1992) developed such models. Wu (1998) describes the asymmetric volatility effects. Again, as above with the leverage models, a key problem with these arguments is that while they can generate asymmetric smiles, they depend crucially on the news event as a trigger. Thus, the models cannot explain price changes, which occur without accompanying news announcements. However, Brown and Jackwerth (2000) find that news about macroeconomic events have very little impact on the smile. Finally, Poterba and Summers (1986) argue that there is only a rather small quantitative effect of the volatility feedback as shocks to market volatility are not very persistent.

Behavioral finance and rational bubbles: The models of behavioral finance argue that a sudden shift in investor sentiment manifests itself in a market crash [Shiller (1989), p.1]. However, it is not clear why such sudden shift in sentiment should be more likely to the downside than to the upside – as long as we do not view fear as an inherently stronger feeling than euphoria. Alternatively, rational bubbles such as the model of Blanchard and Watson (1982) achieve market crashes that are inherently asymmetric. Alas, Gartner (1999) shows that for a rational bubble to start out of equilibrium, individuals would have to form "inconsistent" forecasts – an argument which leads us back towards the behavioral models. Empirically however, these models have not been very successful as the work of West (1988) and Flood and Hodrick (1990) shows.

Stochastic Process: A large literature investigates general stochastic processes for the underlying in order to explain the index option smile. There are deterministic volatility models which keep the framework of a one-dimensional Brownian motion driving the underlying asset price. These models (e.g. Rubinstein (1994) and others surveyed in Jackwerth (1999)) allow the volatility to be a general function of the index function and time. However, if we are willing to assume continuous trading, then the pricing kernel of a one-dimensional process is uniquely determined through Girsanovs Theorem.[12] We explored this avenue by specifying more general functions for the drift and the volatility parameter. Interestingly, when we plot the resulting pricing kernels across wealth, we only add noise: the overall shape stays the same (monotonically decreasing across wealth), and we could not generate the characteristic hump of Figure 2.
Alternatively, we have stochastic models (e.g. Bates (2000) and Bakshi, Cao, Chen (1997)), which add (at least) a second stochastic factor, such as volatility, interest rates, or jumps. Jumps however only give you the negative skewness of the returns (if the jumps are asymmetric) but not necessarily the negative correlation stock vs. volatility. One of the leading proponents of this idea is the stochastic volatility and stochastic jump model of Pan (1999). We simulate her model

[12] Moreover, if trading is not continuous but frequent, one would hope that the associated set of feasible pricing kernels is close to the unique pricing kernel of the continuous trading case.

according to her estimated parameters, but the model-implied pricing kernel is monotonically decreasing. Hence, Pan (1999) cannot explain Figure 2.
One general problem with all these approaches is that the process for the underlying did not change dramatically around the crash. Thus, it is hard to see why we should use a simple Black-Scholes framework precrash and a more complex model postcrash. Alas, one might have wanted to use the more complex model in the precrash period but the option market ignored this fact until the crash, when it realized its mistake. Similar arguments can be made for option pricing models based on the GARCH model, such as Duan (1995). In addition, Duan (1995) assumes a monotonically decreasing pricing kernel, which is inconsistent with Figure 2.

Market imperfections and demand/supply:[13] 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. Demand and supply would determine the final trading price. Grossman and Zhou (1996) suggest such a model with an exogenous demand for portfolio insurance. Unfortunately, they do not generate a sufficiently steep smile, but we will revisit their model below. 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 at all after the 1987 crash. In addition, Brown and Jackwerth (2000) find that large trades of out-of-the-money puts have no obvious impact on the steepness of the smile.[14] Next, in the spirit of Lee and Ready (1991), we looked at the number of seller versus buyer initiated options trades out of a total of 665,548 trades. The distribution of trades is given in table 1. There are slightly more buys (trades near the ask) than sells (trades near the bid) but the imbalance does not seem to explain the smile pattern, in particular as the distribution of trades is insensitive to the choice of moneyness ranges.

[13] In a further exploratory study, we investigated market frictions in a GMM setting starting with the Euler equation 1 = E[m r]. 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. 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.

[14] 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. However, 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.

Table 1. The distribution of trades relative to the last observed bid/ask spread.

Trade above ask	0.07 %
Trade at the ask	0.28
Trade between ask and midpoint	0.10
Trade at the midpoint	0.18
Trade between bid and midpoint	0.09
Trade at the bid	0.22
Trade above the bid	0.06

Further, while the bid/ask spreads 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.
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.
To explain the steep smile with transactions costs for trading the out-of-the-money options is probably insufficient. 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.

Hedging: 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 falling market. However, this problem only applies to investors who desire to stay fully hedged. However, 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.

Section 3.2 – Additional Hypotheses that Explain the Puzzle

Note that the historical sample estimate q^* is backward looking. On the other hand, p^* is derived from option prices, and is forward looking in that it represents the beliefs investors over the life of the index options (31 days in our work). If investors do not have rational expectations, then p^* and q^* may differ greatly.
One form of irrationality we call crash-o-phobia. 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 p^*, which represents prices for dollars received conditional on a large market decline, is raised upward. Given that p^* always integrates to one, p^* needs to decline in other ranges. If in addition the statistical distribution q^* is unchanged – so that the change in p^* represents an irrational change in beliefs - then the pricing kernel may take on the shape shown in Figure 2. In other words, crash-o-phobia is an

irrational change in investor beliefs that shows up in index option prices and therefore in m*. It is a hypothesis that explains our puzzle.[15]

Given that the puzzling behavior of m* appears in the index option data beginning in early 1988, just after the market crash, the crash-o-phobia hypothesis must be seriously considered. 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. 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. 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.

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. However, Figure 2 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. Within the context in which we work, state-dependent utility includes also the case that the representative investor holds assets, e.g. foreign stocks, in addition to the S&P 500 index.

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 future expected marginal utility divided by the current marginal utility. If utility is dependent on a state variable, and that variable is responsible for the puzzle, then changes in the variable must correlate with the index returns. Furthermore, the relation between the 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 or variables v. We then write

[15] Stocks dropped precipitously in Asian and European markets in October 1987, too, and with the same magnitudes as U.S. stocks. In further research, it would be interesting to compare the pricing kernels in these markets with those in the U.S.

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

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 2, the 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 2.

We depict a hypothetical situation of such a state variable v in Figure 4. As wealth increases, state 2 becomes more likely. Thus, the pricing kernel m* is a weighted average of the two pricing kernels associated with the two states of v. Grossman and Zhou (1996) use a similar model to explain the smile in index options as the result of two classes of investors of whom one class is seeking portfolio insurance. Unfortunately, their model does not provide for strong enough effects, so that their aggregate pricing kernel (corresponding to m*) is still monotonically decreasing across wealth. Similarly, one can think about the following potential explanation: For low (high) levels of wealth, utility is based on low (high) marginal utility function of investor 1 (2). In the center, the weights shift and the resulting marginal utility function from the above Figure 1 is obtained. However, it is not clear why the two investors stay segregated and dominate only their respective domains. The investigation of institutional details to explain such separation is ongoing research.

Figure 4. Hypothetical Pricing Kernels depending on State Variable v. We graph the simplest setting where the 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*.

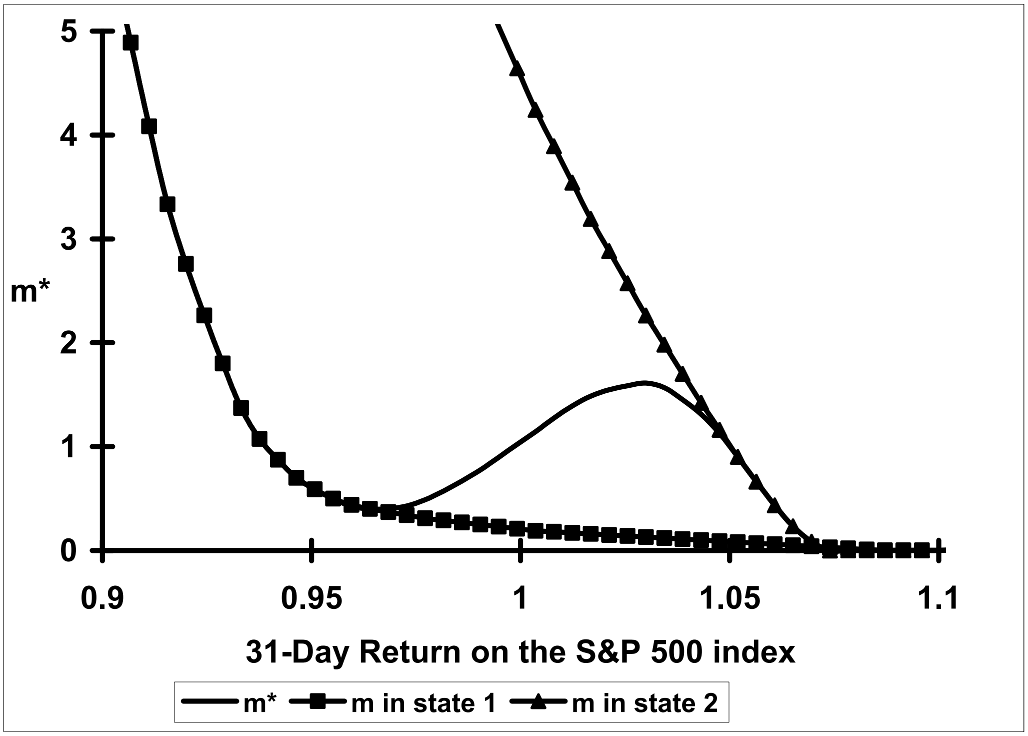

Section 4 - A Model

So far, we have not specified the second state variable. Candidates could be the level of habit persistence or the level of future utility in a recursive utility framework, Epstein and Zin (1989 and 1991). Another class of state variables are (stochastic) volatility, interest rates, or any other asset price which carries additional information about the pricing kernel.

For our model, we 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 as follows. Let the representative agent choose the rate of consumption $C\left(W_t, X_t, t\right)$, and the proportion of wealth in one risky asset $\alpha\left(W_t, X_t, t\right)$, and the proportion in one riskless asset (1-α) to maximize expected utility of lifetime consumption:

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

Where

Eq. 6
$$\begin{aligned} dW_t &= \left(W_t\alpha_t\left(\mu - r_t\right) + W_t r_t - C_t\right)dt + \alpha_t W_t \sigma\left(X_t\right)dz_t \\ dX_t &= -\theta X_t dt + \sigma\left(X_t\right)dz_t \end{aligned}$$

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 each diffusion equation. Changes in the state variable are perfectly correlated with changes in wealth. Using Ito's lemma,

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

Optimal consumption and investment choices satisfy the HJB equation:

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

From the first order conditions:

Eq. 9
$$\begin{aligned} C_t &= e^{-\rho t/B} J_W^{-1/B} \\ \alpha_t &= \frac{-J_W}{J_{WW}W_t}\frac{\mu - r_t}{\sigma^2\left(X_t\right)} - \frac{J_{WX}}{J_{WW}W_t} \end{aligned}$$

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,

Eq. 10
$$J\left(W_t, X_t, t\right) = \mathrm{e}^{-\rho t}\, h\left(X_t\right) W_t^{1-B}/(1-B),$$

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

$$\frac{d}{dC}=0=C^{-B}-h\left(X_t\right)W^{-B}\Leftrightarrow C_t=h\left(X_t\right)^{-1/B}W_t$$

Eq. 11

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

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

Eq. 12
$$r_t=\mu-B\sigma^2\left(X_t\right)+\frac{h'\left(X_t\right)}{h\left(X_t\right)}\sigma^2\left(X_t\right)$$

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

Eq. 13
$$0=\left(1-B\right)\mu h-\rho h+Bh^{1-1/B}-\theta X_t h'+\sigma_t^2\left(X_t\right)\left[\frac{1}{2}h''+\left(1-B\right)h'-\left(1-B\right)Bh/2\right]$$

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:

Eq. 14
$$\frac{m_T}{m_t}=\frac{J_{W_T}}{J_{W_t}}=\mathrm{e}^{-\rho(T-t)}\frac{h\left(X_T\right)}{h\left(X_t\right)}\left(\frac{W_T}{W_t}\right)^{-B}=\mathrm{e}^{-\rho(T-t)}\frac{h\left(X_T\right)}{h\left(X_t\right)}r_{m,t,T}^{-B}$$

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*:

Eq. 15
$$h\left(X\right)=d+\mathrm{e}^{a+bX+cX^2}$$

In order to now empirically test the model, we proceed to an GMM estimation of the model based on a daily time-series (4 April 1986-29 December 1995) of returns of the risk-free rate, the S&P 500 index, and 5 constant 45-day-to-expiration call options. These options range in moneyness from deep-in-the-money (0.96 moneyness), in-the-money (0.98), at-the-money (1), out-of-the-money (1.02), to deep-out-of-the-money (1.04).[16] With the knowledge of the functional form of *h* in Eq. 15, we can then solve for the function $\sigma(X_t)$ as a quadratic form.[17]

[16] Details on the data are in the appendix. A complete description of the dataset is in Buraschi and Jackwerth (2001).
[17] Of the two roots, a positive volatility requires that we take the negative root if (1-B) is less than zero and the positive root if (1-B) is greater than zero.

We can then pick the parameters μ, B, θ, X_0, ρ and the parameters of the h-function, a, b, c, d. Next, we discretize the wealth process and solve for the empirical daily increments in wealth:

Eq. 16
$$dz_t = \frac{r_{m,t,t+\Delta t} - 1 + \left(h(X_t)^{-1/B} - \mu\right)\Delta t}{\sigma(X_t)\sqrt{\Delta t}}$$

where $r_{m,t}$ is the S&P 500 index return from date t to $t+\Delta t$ which proxies for the change in wealth. We can use the increments to update the X-process:

Eq. 17
$$dX_t = \sigma(X_t) dz_t \sqrt{\Delta t} - \theta X_t \Delta t$$

The pricing kernel from date t to $t+\Delta t$ is:

Eq. 18
$$\frac{m_{t+\Delta t}}{m_t} = e^{-\rho\Delta t} \frac{h(X_{t+\Delta t})}{h(X_t)} r_{m,t,t+\Delta t}^{-B}$$

We are now ready to calculate two sets of moments for our GMM estimation. The first is based on the empirical asset returns which are expected to have value 1 when multiplied by m: $E[m_t R_t]-1$ where R_t is the collection of the seven observed asset returns from date t to $t+\Delta t$. The second is based on the first 4 moments of dz_t: $E[dz_t] - 0$, $E[dz_t^2] - 1$, $E[dz_t^3] - 0$, and $E[dz_t^4] - 3$. We collect the moment conditions in a vector g and minimize the GMM objective function $g`\, W\, g$ by varying the parameter vector. Our initial choice for W is simply the identity matrix.
To be continued.

Section 5 – Summary

The goal of out work is to explain the puzzle pictured in Figure 2. 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 investigation is both empirical and theoretical. 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 relation we see in Figure 2. We are estimating a GMM specification of the model and use data on the riskfree rate, the S&P 500 index returns, and on S&P 500 index option returns.
Future research will examine index returns from markets outside the U.S. We will see whether the puzzle is unique to the S&P 500 index options market. The empirical work also will 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:

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

so the derivatives of utility of wealth are

Eq. A2
$$\begin{aligned}
J_W &= \mathrm{e}^{-\rho t}\, h\left(X_t\right) W_t^{-B} \\
J_t &= -\rho\, \mathrm{e}^{-\rho t}\, h\left(X_t\right) W_t^{1-B} / (1-B) \\
J_x &= \mathrm{e}^{-\rho t}\, h'\left(X_t\right) W_t^{1-B} / (1-B) \\
J_{xx} &= \mathrm{e}^{-\rho t}\, h''\left(X_t\right) W_t^{1-B} / (1-B) \\
J_{ww} &= \mathrm{e}^{-\rho t} (-B) h\left(X_t\right) W_t^{-1-B} \\
J_{wx} &= \mathrm{e}^{-\rho t}\, h'\left(X_t\right) W_t^{-B}
\end{aligned}$$

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

Eq. A3
$$\begin{aligned}
0 = {} & e^{-\rho t} \left(h\left(X_t\right)^{-1/B} W_t \right)^{1-B} / (1-B) \\
& + (\mathrm{e}^{-\rho t}\, h\left(X_t\right) W_t^{-B}) \left(W_t \left(\mu - r\right) + W_t r_t - C_t \right) \\
& + (\mathrm{e}^{-\rho t}\, h'\left(X_t\right) W_t^{1-B} / (1-B)) \left(-\theta X_t\right) - \rho\, \mathrm{e}^{-\rho t}\, h\left(X_t\right) W_t^{1-B} / (1-B) \\
& + (\mathrm{e}^{-\rho t} (-B) h\left(X_t\right) W_t^{-1-B}) W_t^2 \sigma^2\left(X_t\right) / 2 \\
& + (\mathrm{e}^{-\rho t}\, h'\left(X_t\right) W_t^{-B}) W_t \sigma^2\left(X_t\right) + 0.5\, \mathrm{e}^{-\rho t}\, \sigma^2\left(X_t\right) h''\left(X_t\right) W_t^{1-B} / (1-B)
\end{aligned}$$

Data

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

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

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. Within each minute we keep the largest set of option quotes which does not violate:

Eq. A5 $$Sd^{-t} \geq C_i \geq \max[0, Sd^{-t} - K_i r^{-t}]$$

We also check for violations of vertical and butterfly spreads.

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.

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

Bernstein, Peter L. “Capital Ideas: The Improbable Origin of Modern Wall Street.” The Free Press (1993).

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

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

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.

Blanchard, Olivier J. and Mark W. Watson (1982). “Bubbles, Rational Expectations, and Financial Markets”, in Paul Wachtel, ed. Crisis in Economic and Financial Structure, Lexington Books, Lexington MA, 295-315.

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

Brown, David P., and J. Jackwerth, “The Information Content of the Index Volatility Smile.“ Working paper (2000), UW Madison.

Brown, David P., and Michael R. Gibbons, “A Simple Econometric Approach for Utility-Based Asset Pricing Models.” Journal of Finance (1985), 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, "ERM Realignment Risk and its Economic Determinants as Reflected in Cross-Rate Options." Economic Journal (1998), 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., and Ludger Hentschel (1992). “No News is Good News: An Asymmetric Model of Changing Volatility in Stock Returns”, Journal of Financial Economics 31, 281-318.

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.

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

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

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

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, "Permanent and Temporary Components of Stock Prices." Journal of Political Economy (1988), 246-273.

Flood, Robert P., and Robert J. Hodrick (1990). "On Testing for Speculative Bubbles", Journal of Economic Perspectives 4, 85-101.

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.

French, Kenneth R., G. William Schwert, and Robert F. Stambaugh (1987). "Expected Stock Returns and Volatility", Journal of Financial Economics 19, 3-29.

Gartner, Manfred (1999). "How Bubbles Are Born and Burst in the Foreign Exchanges", Jahrbücher für Nationalökonomie und Statistik 218, No. 3-4, 294-308.

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, " Equilibrium Analysis of Portfolio Insurance." Journal of Finance (1996), 1379-1403.

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

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

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

Jackwerth, Jens C., and Mark Rubinstein (2001). "Recovering Stochastic Processes from Option Prices", working paper, University of Wisconsin at Madison.

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

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

Lee, Charles M. C., and Mark J. Ready (1991). "Inferring Trade Direction from Intraday Data", Journal of Finance 46, No. 2, 733-746.
Long, D. Michael, and Dennis T. Officer (1997). "The Relation between Option Mispricing and Volume in the Black-Scholes Option Model", Journal of Financial Research 20, No. 1, 1-12.
Lucas, Robert E., Jr., "Asset Prices in an Exchange Economy." Econometrica (1978), 1429-1445
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, "Integrated Time-Series Analysis of Spot and Option Prices." Working Paper (1999), Stanford University.
Pindyck, Robert S. (1984). "Risk, Inflation, and the Stock Market", American Economic Review 74, 334-351.
Poterba, James M. and Lawrence H. Summers, "Mean Reversion in Stock Prices." Journal of Financial Economics (1988), 27-59.
Poterba, James M., and Lawrence H. Summers (1986). "The Persistence of Volatility and Stock Market Fluctuations", American Economic Review 74, 861-880.
Romer, David (1993). "Rational Asset-Price Movements without News", American Economic Review 83, 1112-1130.
Rosenberg, Joshua V., and Robert F. Engle, "Empirical Pricing Kernels." Working Paper (1999), NYU.
Rubinstein, Mark, "Implied Binomial Trees." Journal of Finance, 49, No. 3 (1994), 771-818.
Rubinstein, Mark, "The Valuation of Uncertain Streams and the Pricing of Options." Bell Journal of Economics, (1976), 407-425.
Schwert, G. William (1989). "Why Does the Stock Market Volatility Change over Time?", Journal of Finance 44, 1115-1153.
Shiller, Robert J. (1989). "Market Volatility", MIT Press, Cambridge MA.
Toft, Klaus and B. Prucyk (1997). "Options on Leveraged Equity: Theory and Empirical Tests", Journal of Finance 52, No. 3, 1151-1180.
West, Kenneth D. (1988). "Bubbles, Fads, and Stock Price Volatility Tests: A Partial Evaluation", Journal of Finance 43, 639-656.
Wu, Guojun (1998). "The Determinants of Asymmetric Volatility", working Paper, University of Michigan Business School.