Draft of 8/5/2014

Comments welcome

Options on Leveraged ETFs: A Window on Investor Heterogeneity

Stephen Figlewski* and Muhammad Fahd Malik**

* Professor of Finance	** Quantitative Strategist
New York University Stern School of Business	AllianceBernstein
44 West 4th Street, Suite 9-160	1345 Ave of the Americas
New York, NY 10012-1126	New York, NY 10105
sfiglews@stern.nyu.edu	mfm341@nyu.edu
tel: 212-998-0712	

Options on Leveraged ETFs: A Window on Investor Heterogeneity

Abstract

A risk-neutral probability distribution (RND) for future S&P 500 returns extracted from index options contains investors' true expectations and also their risk preferences. But the empirical pricing kernel that emerges in a representative agent framework, which suppresses investor differences, is inconsistent with investor rationality. The anomalous shape can be generated easily in a model with heterogeneous investors facing limits to arbitrage, however. We explore investor heterogeneity directly via RNDs extracted from options on three exchange traded funds with leveraged short and long exposures to the S&P 500 index, and find large differences that are not arbitraged away. For example, holders of ETFs with short exposure to the market value payoffs in negative return states significantly more than those with long exposure do. Separating true expectations from risk premia requires further assumptions, so we consider polar cases in which either all investors have the same expectations but different risk preferences, or the reverse. The results are largely consistent with the expectations differences we anticipate from investors choosing short or leveraged long market exposure. We then look at changes in RNDs to see how realized returns affect investors with different expectations. A large negative realized return raises the median future return expected by bulls but lowers it for bears. Uncertainty over future returns widens for both types of investors when they are wrong and narrows when they are right.

JEL Classifications: G14, G13, G11

Keywords: investor heterogeneity; risk neutral density; exchange-traded funds

Introduction

In standard asset pricing theory, an investor's demand for any financial asset is a function of his or her beliefs about the probability distribution for the payoff in each possible future state of the world, modulated by a "stochastic discount factor" (SDF) or "pricing kernel" that gives the value today of $1 in each future state. The investor assesses the fair value C for the asset by

$$C = \int_X g(x)\, p(x)\, k(x)\, dx \qquad (1)$$

where g(x) is its payoff in future state x, $x \in X$, where X denotes the state space; p(x) is (the investor's subjective estimate of) the true probability density over the states in X, and k(x) is the pricing kernel that discounts the payoff in each state x.

The market aggregates individual investor demands into an excess demand function that clears the market at the current equilibrium price. At that price, optimistic investors who calculate a high value for C will overweight the security in their portfolios and pessimistic investors will underweight it or even sell it short.

Under fairly general assumptions the aggregation can be thought of as happening at an earlier stage, with heterogeneous individuals being combined into a "representative investor" who holds average beliefs and average risk preferences. Often the representative investor is simply called "the market." The utility function and probability beliefs of this generic investor generate the SDF. In this sense, it is reasonable to talk about "the market's" expectations and risk preferences.

The market evaluates a contingent dollar payoff in each possible future state of the world and assigns a price to it, which reflects the value of a dollar received in that state (a "state claim") times the probability of that state occurring. The price of a state claim is frequently called the Arrow-Debreu price, and the set of Arrow-Debreu prices over the full state space is the pricing kernel. Any security can then be valued easily: For each future state of the world, multiply the dollar payoff in that state by its Arrow-Debreu price and add them up across all states. Any violation of this pricing relationship creates a profitable arbitrage trade.

The pricing kernel embodies the market's risk preferences and beliefs about the true or objective probability density over future states, p(x), which is commonly referred to as the p-density. Harrison and Kreps (1979) proved a deep and important result: For a market with no arbitrage opportunities, the market's beliefs about p(x) and any risk premium embedded in the pricing kernel k(x), can be combined and represented by a modified probability distribution with the property that a risk neutral investor facing this "risk neutral density" (RND) would value every state-contingent payoff exactly the same way as it is priced in the actual market, under the true density p(x) with pricing kernel k(x). The RND is often called the q-density and we will use the terms interchangeably.

That is,

$$q(x) = e^{r_f T}\, p(x)\, k(x) \tag{2}$$

where r_f is the riskless interest rate. The $e^{r_f T}$ term converts the pricing kernel, which gives the present value of $1 received in state x, into its equivalent in terms of a sure $1 received in the future T periods later.

Ross (1976) demonstrated how options can expand, and potentially complete, the market for payoffs that are contingent only on the price of some underlying asset. Breeden and Litzenberger (1978) then showed how the risk neutral density q(x) can be extracted from the prices of a set of options with a continuum of strikes and a common expiration date T. If C(x,T) is the current market price of a call option with exercise price x and maturity T, then

$$q(x) = \; e^{r_f T}\, \frac{\partial^2 C(x,T)}{\partial x^2} \tag{3}$$

The RND for a future stock price x is equal to the second partial derivative of the call value function with respect to the strike price future-valued to the expiration date T. The surprisingly simple derivation of this important result is shown in the Appendix.

Given a q-density extracted from option prices and an estimate of the true probability density p(x), the pricing kernel can be computed as

$$k(x) = \; e^{-r_f T} q(x)/p(x) \tag{4}$$

With these theoretical results and the development of an active market in stock index options, much research has been done to exploit and explore these relationships.[1]

In a market with a range of traded options, the q-density can be observed, or more precisely, it can be closely approximated. One would like to be able to use the q-density to learn about investors' true probability expectations and/or their risk preferences. Attempts to extract probability beliefs from the RND include Jackwerth and Rubinstein (1996), who looked at the S&P 500 index around the Crash of 1987; Gemmill and Saflekos (2000), who tried to learn about the market's expectations for the British election of 1997; Melick and Thomas (1997), who analyzed the behavior of the RND at the time of the first Iraq war in 1991; Breeden and Litzenberger (2013), who used cap and floor prices to assess expectations about the future path of interest rates; and many others. Both the U.S. Federal Reserve and the Bank of England compute RNDs from different markets in evaluating economic conditions.[2] Papers focusing on using RNDs to explore investor risk preferences include Bliss and Panigirtzoglou (2004), Jackwerth (2004), Bates(2000), Liu et al (2007), and others.

[1] See reviews of the literature on risk neutral densities by Jackwerth (2004) or Figlewski (2010).
[2] See Federal Reserve Bank of Minneapolis, (2014) and de Vincent-Humphreys and Noss, (2012).

Although these efforts are suggestive, in general it is not possible to separate expectations from risk preferences without additional assumptions to constrain the problem. Some authors propose either a specific form for the p-density that allows the kernel to be computed, such as lognormal with fixed mean and standard deviation, or for the short run returns process, such as GARCH.[3] Others, like Bliss and Panigirtzoglou (2004) model the representative agent's utility function as being of a certain form, typically with constant relative (CRRA) or absolute (CARA) risk aversion, and proceed to extract an implied p-density. It is possible to extract both the market's p- and q-densities with a large enough set of repeated observations on panels of option prices over time, as long as both are assumed to be constant in the data sample.[4]

Much of the work in this area has focused on the S&P 500 index. This index is widely chosen as a proxy for the stock market portfolio, and in turn as a proxy for total wealth. In that case, higher stock prices correspond to greater overall wealth, and if the representative investor has a normal utility function, the marginal value of $1 goes down when wealth increases, producing a monotonically decreasing pricing kernel over S&P 500 index returns.

We have extracted RNDs from options on the SPDR exchange traded fund (ETF), often known by its ticker symbol SPY. The SPY is a share in a stock portfolio designed to track the performance of the S&P 500 index. It is an investment vehicle that allows a small investor to take a long position in the whole market portfolio in the same way that she might buy a share in a single company. Options on the SPY with an extensive set of exercise prices and maturities are actively traded.

We have data on SPY call and put options between 2007 and 2014, for a sample of 1440 trading days. To estimate p-densities for those dates, we assume that investors believe the density is lognormal with a mean equal to the current riskless rate plus a risk premium of 5.00%.[5] For the market's estimate of the volatility through option expiration, we combine realized volatility over the previous three months and the current VIX implied volatility index, using weights that would have given the most accurate predictions of realized volatility over comparable horizons in the past. In this way, investors are assumed to treat the mean and standard deviation as being fixed over the life of the option, but changing over longer horizons. Although the density is conditionally lognormal, the empirical distribution from longer samples will have fat tails. Details on the sample selection and data handling procedures are provided in the Appendix.

---

[3] Jackwerth and Rubinstein (1996) fit a binomial model for the S&P 500 but need to constrain the risk neutral density to be lognormal. Rosenberg and Engle (2002) fit a GARCH model to returns. Gemmill and Saflekos (2000) use a mixture of two lognormals to reflect two possible outcomes of the British election of 1997. Rompolis and Tzavalis (2008) suggest the Gram-Charlier density approximation technique, while Eriksson et al (2009) prefer the Normal Inverse Gaussian, and many other alternatives have been explored in the literature.

[4] See Ait-Sahalia and Lo (1998, 2000) for an empirical demonstration and Ross (2014) for a formal proof of this proposition.

[5] This risk premium was chosen based on surveys reported by Fernandez and del Campo (2010) who found that for analysts and companies in the U.S. and Canada, the median expected risk premium over the riskless interest rate was 5.0% in April 2010. Other available surveys give roughly comparable values.

Figure 1 plots the pricing kernel for May 2, 2008 extracted from one-month options on the SPY ETF. The vertical dashed line is at the current level of the S&P index implied from the SPY. The left portion of the curve shows how the representative investor would value an additional $1 in the range of S&P index returns representing a substantial loss relative to initial wealth. It displays the expected negative relationship between the future index level and the kernel. But there is an anomalous hump in the region of gains. On this date, the market was willing to pay more for a $1 payoff if the stock index at option expiration was 100 points above today's level than if it was 100 points lower. This pattern is not uncommon. Indeed, the hump was present for most of our sample days and it has been observed by many researchers in the past.[6] It has been termed the "pricing kernel puzzle."

[insert Figure 1 about here]

One very promising solution to the puzzle has been proposed by Shefrin (2008), who shows it can arise naturally from heterogeneity among investors. The standard representative agent assumption allows enormous simplification of pricing models, but it has difficulty with zero net investment contracts like options, since identical investors would not take positions on opposite sides of the same contract. In Shefrin's model, there are two classes of investors, optimists and pessimists. They have identical CRRA utility functions but differ in their beliefs about both the mean and the standard deviation of the returns distribution. The optimists expect higher mean returns, so they want more exposure to the market on the upside. If they also expect lower volatility than the pessimists do, the resulting shape of the pricing kernel can look like Figure 1.[7]

Our goal in this paper is not to explain the pricing kernel puzzle, but rather to take its existence as evidence of an important feature of our financial markets that cannot easily be explained within our standard representative investor models. Real world investors are heterogeneous in returns expectations, risk preferences, initial portfolios, and every other important dimension. There are also limits to arbitrage and other frictions that allow active trading in many zero net investment derivatives, even closely related instruments whose payoffs are based on the same underlying asset price. These are redundant in a frictionless theoretical market, but may attract clienteles of specific subsets of real world investors. We hope to learn about investor heterogeneity by examining RNDs extracted from exchange traded funds that provide different leveraged exposure to the S&P 500 index.

The market for ETFs has expanded rapidly and now includes multiple contracts based on the same underlying index, but differing in the direction and leverage of the exposure. In this paper, we focus on SPY and two other such ETFs, that will be identified by their ticker symbols: SSO,

[6] See Hens and Reichlin (2013) or Brown and Jackworth (2012), for example.
[7] In a later paper, Barone-Adesi et al (2013) develop a somewhat different sentiment-based model with investors who are both over-optimistic and over-confident, that is also capable of generating a pricing kernel with an anomalous upward sloping region.

which tries to produce twice the daily returns on the index, and SDS, an "inverse" ETF that is designed to provide twice the daily return of a short sale of the S&P index. For example, if the S&P return on some date t is +1.0%, returns on the SPY, the SSO, and the SDS should be, respectively, +1.0%, +2.0%, and -2.0%.

All three ETFs have traded options from which RNDs can be extracted. In a completely frictionless market, arbitrage among the ETFs and the underlying portfolio of S&P 500 stocks should lead to perfectly consistent pricing, such that all of the derivatives would, in fact, be redundant. But such arbitrage does not eliminate all pricing differences among these ETFs in the real world. The flourishing trade in all three suggests that each attracts its own clientele of investors, who can be expected to differ in their probability beliefs and/or their risk attitudes in significant ways from those who are attracted to the other contracts.

Computing implied volatilities from option prices is a long-established procedure. Extracting the entire RND gives much greater insight into how the market values payoffs under different future states of the world and does not require assuming the validity of a specific option pricing model. But we are not able to separate the resulting RNDs into expectations and risk preferences, nor is there any easy way to explore heterogeneity within the investor population, even though active trading in zero net supply contracts would not exist without it. One of the main contributions of this paper is to demonstrate the consistent and highly significant differences in the risk neutral probabilities over the same set of contingencies that are produced by investor populations who are expected to have quite different expectations and perhaps risk preferences, as well. For example, as one would expect, RNDs extracted from the double short SDS contracts have lower medians than those from SPY and SSO. We introduce a new tool, the Relative Demand Intensity (RDI), to analyze how two RNDs are related to one another over different regions of the returns space. The RDIs show that investors on the short side assign lower risk neutral probabilities to a future rise in the market than do longs, but both SPY and SSO investors value payoffs in negative states in which they would have severe losses more highly than the shorts do.

Without separating RNDs into expectations and risk preferences, we are able to generate interesting results by considering alternative polar cases that weaken the homogeneous investor model: either investors in all three ETFs have the same expectations about the p-density but differ in their risk tolerance, or they all have the same risk aversion but different probability beliefs.

Assuming identical p-density expectations produces three different pricing kernels for each day. If investors only differ in risk aversion, SSO holders are the most risk tolerant, then SPY, and finally SDS investors who use contracts with short exposure as a hedge. However, the results we find under this assumption are clearly inconsistent with this ordering. Assuming equal risk tolerance instead, we can imply out different probability beliefs from the RNDs. The relationship between the p-densities for the long exposure contracts versus those for SDS are as anticipated, with the SSO and SPY investors assigning higher probabilities to up markets and

SDS investor higher probabilities to future market declines. However, the relationship between the single long SPY and the double long SSO p-densities is muddier. This is one of several places where the empirical results show the differences between the two long exposure ETFs are more ambiguous than anticipated.

Moving beyond static comparisons of RNDs and the curves we extract from them, we then look at how each investors in each set respond to the same price shock contained in the most recent market return. Among the interesting results, we find that long exposure investors raise (lower) their risk neutral expectations for future index returns following a positive (negative) realized return on date t, but SDS investors do the opposite. Having probability distributions, not just single parameters like an implied volatility, lets us explore how higher moments behave. We find, for example, that when investors are wrong, as when an SDS holder sees a sharp increase in the index, the interquartile range (as a nonparametric alternative to standard deviation) of their RND widens and it narrows when they are right.

The remainder of the paper proceeds as follows. In Section 2, we briefly describe the data sample and our procedures for extracting risk neutral densities and converting them to a common base defined in terms of a standardized p-density so they can be compared. The full details of these technical issues are presented in the Appendix. In Section 3 we explore how the RNDs from the three ETFs differ when observed on the same date. Section 4 looks at how each of them responds to the most recent realized return in the market. Section 5 concludes.

2. Data & Methodology

End-of-day market bid and ask prices for SPY, SDS and SSO options come from OptionMetrics. The two levered ETFs began trading in the middle of 2006, but options on them were only introduced towards the end of 2007. Our sample starts in October 2007 and ends in March 2014. This gives us 1615 observations for SPY, 1614 observations for SSO and 1611 for SDS. After eliminating dates without adequate data to compute RNDs on all three ETFs, and removing days with obviously erroneous probability densities (mainly due to violation of no-arbitrage conditions in the reported prices) we have 1440 usable densities for each of the three ETFs.

The ETF options' payoffs are based on expiration day prices for their underlying ETFs. The Appendix describes in detail our procedure for extracting the risk neutral densities over future ETF values from these options. To make them comparable between ETFs and also across RNDs from options with different time to expiration, we convert first into RNDs defined on the common support of S&P 500 index levels, and then into standardized returns expressed as standard deviations around the expected mean S&P index return. In some of the analysis, we report average differences between curves over subintervals of the space of standardized returns. Table 1 shows the number of days with available observations for each ETF at points along a range from -3.0 to +3.0 standard deviations. Because the stock market rose sharply over much of the sample period, there is very good coverage only in the intervals between -3.0 and +1.5, and even +1.5$\hat{\sigma}$ is a little thin for SDS.

[insert Table 1 about here]

Figure 2 shows the transformed RNDs for May 2, 2008 extracted from June options with T = 50 days to expiration. On this date, the S&P 500 index closed at 1413.90, up 0.32% from the previous day. The VIX index was at a relatively moderate level of 18.18, down 0.70 from the day before, and the blended volatility used to construct the p-density and to convert RNDs on prices to their equivalents in returns was 15.75.

[insert Figure 2 about here]

The figure reveals several common features of the data. The curves are roughly similar, but the q_{SPY} density is available over a broader range than the others. Also, all of the curves are truncated at the right end (higher S&P levels). One problem in the analysis is that an RND can only be computed from the data in the range spanned by the option strike prices available in the market. After a large move in the underlying index, new strikes are added to widen the range, but the new contracts may not extend very far into the tail of the density, especially for the less active leveraged contracts.

The inability to fit the tails from market data has led researchers to a variety of solutions to complete the densities.[8] Without a complete density, it is not possible to compute the mean, standard deviation, or other moments. However, it is possible to compute quantiles in the observed portion of the density and also the total probability in each tail. Rather than depending on assumptions about the shape of the RND tails, in this work we focus on nonparametric statistics: the median, rather than the mean, and the interquartile spread in place of the standard deviation. Many of our tests involve comparing average probabilities over different ranges of the observed portions of the curves.

[8] Fitting a parametric density to the data imposes that density's shape on the tails by assumption. Birru and Figlewski (2012) suggest extracting the middle portion of the RND nonparametrically from the options market and completing it with tails constructed from the Generalized Pareto distribution. Bliss and Panigirtzoglou (2002, 2004) and others also use nonparametric estimation but effectively constrain the tails to be lognormal by assuming Black-Scholes implied volatility is constant beyond the range of strikes available in the market.

3. Comparative Statics of Risk Neutral Densities and Related Functions

Investors in these ETFs are entering into contracts that specify very different payoffs as a function of the realized future value of the same variable. It is reasonable to expect these investors to differ in their beliefs about the true p-density and possibly in their risk attitudes, as well. For example, an investor who buys a call option on the double short SDS ETF probably expects the S&P 500 index to be lower in the future than does a buyer of a call on the double long SSO. Or he could be more averse to the risk of a fall in the index and is hedging by buying an option that will pay off in that unfavorable state of the world, while a more risk-tolerant SSO investor is increasing her exposure to stock market risk at the same time. In this section we will propose and test a number of hypotheses about how the RNDs from the three ETFs relate to one another.

Any discrepancy in the risk neutral densities extracted from derivative securities based on the same underlying asset indicates an arbitrage opportunity. In a frictionless financial market, all such differences should be arbitraged away. The first hypothesis, H0, is that there are no significant differences among the RNDs derived from options on the three ETFs. As will be immediately clear, many such differences do exist and are highly statistically significant. We will be able to conclude *a fortiori* that the hypothesis that arbitrage eliminates any discrepancies among the risk neutral densities in these markets is refuted.

The Medians of the Risk Neutral Densities

Since persistent RND differences do exist, the above reasoning suggests several hypotheses about their relative positions in the three markets.

H1a: The median of the q_{SSO}-density will be higher than the median of the q_{SPY}-density.

H1b: The median of the q_{SSO}-density will be higher than the median of the q_{SDS}-density.

H1c: The median of the q_{SDS}-density will be lower than the median of the q_{SPY}-density.

The top portion of Table 2 shows the average and standard deviation of the RND medians from each of the three markets over the 1440 days in our sample. The median of the SPY q-density is a little above a return equal to the riskless rate (which would be 0.0 standard deviations here). The mean of the market's risk neutral density should equal the riskless rate, and if the density is left-skewed, as it is for the S&P 500, the median will be above the mean. The SSO median is also positive but averages a little below the SPY median, while the average SDS median return is below the riskless rate (with a substantially higher standard deviation than for the other ETFs). The lower portion of the table provides more insight into these differences in medians.

[insert Table 2 about here]

The average of the medians over time can be influenced by days that rank as outliers, so we first simply report the fraction of days for which the median of the first listed ETF minus the median of the second was positive. Here the median of the double long SSO's risk neutral density was above that of the SPY on only about 30% of the sample days, but medians for both ETFs with long exposure to the S&P index were above the median of the double short SDS more than 97% of the time, by more than one third of a standard deviation (of standardized return) on average. The standard deviation of the difference in medians across days was only 0.042 for SSO-SPY, but given the sample size, the average difference was still highly significant. The differences between SSO and SPY versus SDS are much larger on average and highly significant but much more volatile.

Table 2 strongly supports hypotheses H1b and H1c that the SPY and SSO medians are above the median of the double short SDS. But the double long SSO's median is less than that of the SPY, contradicting H1a. This interesting result illustrates the interplay between investors' expectations and risk aversion. SSO buyers very likely believe the expected return on the S&P index is higher than SPY buyers think, but a given percentage fall in the S&P index produces losses twice as large for SSO than SPY investors, which suggests they might also place greater value on downside protection.

The RNDs for May 2, 2008, plotted in Figure 2 are consistent with this line of reasoning. In the region of gains, the SSO q-density is mostly above that for the SPY (and both are well above the SDS RND). In the region of small losses that are less than one standard deviation below the mean, SSO is below SPY, suggesting that the double long investors think an outcome in this range is less probable than SPY investors do. But larger losses of two or three standard deviations in the index are much more painful for an SSO investor than an SPY investor and should cause him to value a payoff in that state of the world more highly, which is what we see in Figure 2. SDS investors weight payoffs in a falling market more highly than do investors in either of the long exposure ETFs.

Relative Demand Intensity across the Returns Space

To explore these issues in greater detail, we can compare average risk neutral probabilities over different regions of the returns space. If buyers of the double long SSO are at least partly doing so because they anticipate higher returns on the market portfolio than the average investor does, we expect $q_{SSO} > q_{SPY}$ in the region of gains and $q_{SSO} < q_{SPY}$ in the region of (moderate) losses, and similar reasoning leads to equivalent hypotheses about the other comparisons among these RNDs. But in the presence of risk aversion, the different exposures to stock market risk these contracts have may confound this pattern based on expectations.

The ETFs we are examining are all tied to the S&P 500 Index, which is frequently taken to be a proxy for the representative investor's total wealth, as in discussions of the equity premium puzzle. A nonmonotonic pricing kernel is a puzzle because the relationship between the index and total wealth is expected to be monotonically increasing. But ETFs and other securities with leveraged exposures now exist for many underliers that cannot be considered as proxies for total wealth, including indexes covering narrower sectors of the stock market, as well as those based on other asset classes.[9] We can explore relative intensities of demand for exposure to different portions of the returns space for these underlying assets by comparing the RNDs extracted from options on their leveraged and inverse ETFs.

For two ETFs, A and B, with different exposures to the same underlier, define the two sets of investors' Relative Demand Intensity (RDI) for a payoff at a given return x by the ratio of their RNDs at x:

$$RDI(x) = q_A(x)/q_B(x) \tag{5}$$

The Relative Demand Intensity shows how A and B investors differ in the way they value payoffs under the same contingency. If they put the same value on a payoff when the standardized return on the underlying asset is x, then RDI(x) = 1.0.

RDIs vary widely over time, even from one day to the next. Figure 3 shows averages over the entire sample for three RDIs. The dotted line is the relative demand intensity for the double short SDS relative to the unlevered SPY. That is, A and B in equation (5) represent SDS and SPY, respectively. RDI(x) above 1.0 indicates that SDS investors have stronger demand than SPY investors do for a payoff at a future S&P 500 level corresponding to a standardized return of x. The interplay between expectations and risk tolerance can be seen clearly here.

[insert Figure 3 about here]

On the right, in the region of small to medium-sized gains on the S&P 500, SDS investors are less eager than SPY investors for exposure. They consider a positive excess return on the index to be less likely than SPY investors expect. But the losses SDS holders would experience with a big rise in the index should cause them to want to hedge against that event, and we see that the relative value they place on a payoff if the market rises begins to increase for returns above about 2/3 of a standard deviation. At the far right end of the graph, the double short investors may become more eager for a contingent payoff than are SPY investors with long exposure. We have

[9] Examples that are currently available in the market include leveraged ETFs based on the Dow Jones, the Russell 2000, and the NASDAQ 100 stock indexes, as well as on bonds, commodities, currencies and even real estate.

very few data points at standardized returns above 1.5, but those few do show RDIs for SDS relative to SPY that are well over 1.0.

The reverse pattern holds on the downside. SDS investors expect the market to fall, so they value payoffs in that region consistently higher than do SPY investors, until the losses are worse than -2.0 standard deviations. At that point, SPY holders' aversion to very large losses leads them to want insurance, while SDS holders would be comfortably experiencing a substantial gain.

By contrast, the dashed line displays the relative demand intensities for SSO double long investors relative to SPY investors who hold unlevered long exposure to the index. The pattern here is weaker than when comparing short and long investors. On the downside, the leveraged buyers are ready to pay more at all return levels than unlevered ones, consistent with aversion to the double-sized losses they will experience for any given negative x. There is very little difference between the two in the region of small gains, but the graph suggests that SSO investors are a little more eager for payoffs for returns above one standard deviation.

Finally, the solid line compares the preferences of SSO versus SDS investors. For payoffs on moderate-sized returns we see strong relative demand differences in the expected direction, with SDS holders liking the downside and SSO holders liking the upside. But interestingly, for extreme returns at both ends, these preferences come together, consistent with the idea that SSO investors think a very large drop in the index is less likely than SDS investors do, but they will pay more for exposure to that contingency in order to mitigate the big losses they would sustain if it does happen. The situation is reversed on the upside, where SDS investors would consider a moderate rise in the S&P to be less probable than SSO investors expect, but are ready to buy insurance against large losses in case of a strong rally.

Because of data limitations, as shown in Table 1, we cannot examine the full range of x-axis values. We therefore focus on three regions of the S&P 500 return space, defined in terms of the current volatility forecast $\hat{\sigma}$: small to medium gains (0.0 to +1.5 $\hat{\sigma}$), small to medium losses (-1.5 to 0.0 $\hat{\sigma}$), and large losses (below -1.5 $\hat{\sigma}$). The general result we hypothesize is that expectations will dominate in the region of small gains and losses on the index, but risk preferences will dominate in the region of large losses.

With three ETFs and three regions of returns, there are nine basic hypotheses.

The specific hypotheses H2a - H2i are as follows.

In the region of moderate gains (0.0 to +1.5 $\hat{\sigma}$),

H2a: RDI(x) > 1.0 for SSO/SPY;
H2b: RDI(x) < 1.0 for SDS/SPY;
H2c: RDI(x) > 1.0 for SSO/SDS.

In the region of moderate losses (-1.5 to 0.0 $\hat{\sigma}$),

H2d: RDI(x) < 1.0 for SSO/SPY;
H2e: RDI(x) > 1.0 for SDS/SPY;
H2f: RDI(x) < 1.0 for SSO/SDS.

In the region of large losses (below -1.5 $\hat{\sigma}$),
H2g: RDI(x) > 1.0 for SSO/SPY;
H2h: RDI(x) < 1.0 for SDS/SPY;
H2i: RDI(x) > 1.0 for SSO/SDS.

To be precise about the test statistics, the test of hypothesis H2a is as follows. We want to look at the average value of RDI(x), defined as $q_{SSO}(x) / q_{SPY}(x)$, over the range of standardized returns from 0.0 to +1.5. For each date in the sample, RDI(x) is averaged over all x values in the range for which we have been able to extract risk neutral densities from the available options data. In cases where a curve does not extend all the way to the end of the region, we compute the average differences over the available portions. From the daily averages of RDIs in this interval we compute the fraction of days with RDIs above 1.0, their mean and standard deviation, and the t-statistic on the difference between the mean and 1.0. Table 3 shows the results.

[insert Table 3 about here]

The results in Table 3 provide strong support for six out of the nine hypotheses. H2a held that SSO investors would be more eager than unlevered SPY investors for exposure in the region of moderate gains. Although this was true on 50.0% of the sample days, the average RDI(x) for SSO/SPY was 0.997, significantly below 1.0 in a one-tailed test, so H2a is significantly rejected. This could mean that SPY investors had more optimistic expectations for returns on the market than SSO buyers had. But another possible explanation is that SSO buyers' expectations were much more positive than those of SPY investors, so they may have estimated the probability of only a small positive return as being lower than what SPY investors expected.

Similarly, hypothesis H2d that SSO investors would have lower demand relative to SPY investors for a payoff in the region of moderate losses was also refuted, more decisively than was H2a. If SSO investors are more optimistic than SPY investors, the point at which expectations differences (SSO investors predict a smaller probability of a down market than SPY investors anticipate) are outweighed by risk aversion (SSO investors are hurt more by a drop of a given size) appears to occur at smaller losses than the region covered by H2d.

H2h, the hypothesis that risk preferences would lead SPY investors to value payoffs for large down moves more highly than SDS investors do, was also refuted by the test on the average, even though H2h was true on 50% of the days. Figure 3 suggests that this result could reflect higher average relative demand intensity for SDS over SPY for moderately large losses

($-2.0\,\hat{\sigma} < x < -1.5\,\hat{\sigma}$) but stronger demand from SPY investors in the region of very large losses ($x < -2.0\,\hat{\sigma}$).

Pricing Kernels

We would like to be able to separate the differences in these RNDs into the part due to different beliefs about the true p-density and the part that reflects different risk attitudes. This is not possible without additional assumptions, but we are able to explore two polar cases: either all investors have the same beliefs about the true returns distribution but different risk preferences, or alternatively, investors all have the same utility functions and only differ in their expectations about the returns distribution. We will examine these two cases in turn.

In our setup, the pricing kernel k(x) gives the value to the investor of $1 received in T days if the stock price at that date S_{t+T} is x standard deviations away from the mean return, relative to the present value of $1 received for certain at date t+T. If investors agree on the p-density but differ in how they risk-neutralize it, their different RNDs will produce different pricing kernels, which we will denote as $k_{SPY}(x)$, $k_{SSO}(x)$, and $k_{SDS}(x)$.

We have computed pricing kernels from the three RNDs shown in Figure 2, using the same p-density for each. We assume investors agree that the p-density for the S&P 500 index is lognormal with a log mean equal to the current annualized riskless interest rate $r_{f,t}$ less the index dividend yield d_t plus a risk premium of 5.0%, all scaled appropriately for a T-day horizon. The annualized volatility input is calculated using equation (A18).

Using the q-densities extracted from the options market and the p-densities constructed as just described, we computed pricing kernels for our 1440 day sample. Figure 4 displays these pricing kernels for the single day May 2, 2008. Note that they all exhibit the anomalous nonmonotonic section in the region of moderate returns. If the explanation for this shape is investor heterogeneity, Figure 4 shows that each market contains investors with diverse beliefs. This is not really surprising, given that every option contract has to have a buyer and a seller.

[insert Figure 4 about here]

If all investors agree on the true density, differences in the kernels must be due to different risk preferences. Consider an investor who takes a long position in the double long SSO versus one who buys the unlevered SPY ETF. For a given drop in the underlying index the loss to the SSO is twice that on the SPY, and on the upside, the SSO gains twice as much as the SPY. Given identical probability beliefs, investors who choose the SSO would likely be less risk averse than SPY investors.[10] Relative to SPY buyers, a less risk averse investor has less demand for a hedge against a given drop in the market, and more for exposure on the upside, since the marginal utility of wealth for a less risk averse investor falls off more slowly as wealth increases.

The effects on the pricing kernels might be most visible for small and medium-sized moves, particularly on the downside where an SSO buyer could sustain a serious hit to wealth from a sharp drop in the market. If we look far enough into the left tail, we would expect to find $k_{SSO}(x)$ rising above $k_{SPY}(x)$ for extreme losses. Figure 4 shows that on this date, this transition happened at a loss level of around -1.0 standard deviations.

Yet, these comparisons are ambiguous because a given move in the index produces different dollar returns to the two ETFs. The loss an SPY holder experiences when the market return is 1.0 standard deviations below expectations is what an SSO holder gets with an index loss half that size. If an SSO buyer is only a little less risk averse than an SPY investor, he will be less concerned about a dollar loss of a given size, yet he might well want to pay more to mitigate the utility loss from a given negative S&P 500 return, which produces twice as bad a loss for him in dollars as for an SPY investor. The issue is how quickly the SSO buyer's aversion to losses increases as they become larger, and similarly, how quickly his marginal utility of wealth falls as profits increase on the upside. Given the offsetting influences, we will report the comparisons between $k_{SSO}(x)$ and $k_{SPY}(x)$ below, but will not try to specify formal hypotheses about how they should be related.

Now consider $k_{SDS}(x)$. It is reasonable to assume that regardless of what position is taken in these particular ETFs, nearly all investors have net positive exposure to the stock market overall. With identical expectations as SPY buyers about the p-density, an investor will only choose to go long in the double short SDS ETF as a hedge, because he is more risk averse than average. He will pay more than an SPY or SSO investor to receive a payoff when the market goes down, and less for a payoff when the market is up.

This reasoning leads to six hypotheses about the relative positions of the pricing kernels extracted assuming all investors agree on the true p-density within the three regions of the returns distribution we examined above. SSO investors are the least risk averse, next the SPY investors and then those who go long the double short SDS.

The specific hypotheses H3a - H3f are as follows.

---

[10] They may also be more liquidity-constrained and willing to pay more for leverage. For either of these to affect pricing, there must be limits to arbitrage that create some segmentation between the two markets, so that, for example, a risk averse SPY investor would not simply create the identical exposure to the index by taking a position half as large in SSO. The evidence we present in this paper clearly demonstrates that the markets for options on these closely related ETFs are indeed somewhat segmented, particularly between the two ETFs with long exposure versus the double short SDS. However, we also think the main difference between the clienteles for these instruments is more likely to be in their returns expectations than their risk preferences.

In the region of moderate gains (0.0 to +1.5 $\hat{\sigma}$),

H3a: $k_{SDS}(x) < k_{SPY}(x)$;

H3b: $k_{SSO}(x) > k_{SDS}(x)$.

In the region of moderate losses (-1.5 to 0.0 $\hat{\sigma}$),

H3c: $k_{SDS}(x) > k_{SPY}(x)$;

H3d: $k_{SSO}(x) < k_{SDS}(x)$.

In the region of large losses (below -1.5 $\hat{\sigma}$),

H3e: $k_{SDS}(x) > k_{SPY}(x)$;

H3f: $k_{SSO}(x) < k_{SDS}(x)$.

.

[Table 4 about here]

Table 4 contains results relevant to the six hypotheses listed above. H3a and H3b, covering the region of moderate gains, are strongly confirmed, with the kernel from SDS lying well below those from SSO and SPY more than 70% of the time, by an average amount that is significant with a t-statistic around 30. In the region of moderate losses on the S&P index, the SDS kernel is above the SPY kernel as hypothesized (H3c), but below the kernel from SSO, which should not happen if the only difference between the two sets of investors were risk aversion. The supposedly less risk averse SSO investors also produce a kernel over negative returns significantly above that from SPY on average. Finally, for extreme losses both of the long-exposure kernels are above the SDS kernel, refuting both H3e and H3f. These results call into question the joint hypothesis that is embedded in H3a - H3f, which is that the investors in the three ETFs project identical p-densities and only differ in risk tolerance.

Implied p-densities

It is not really plausible that these leveraged and inverse ETFs were developed to cater to investors who have identical expectations about returns on the S&P 500 index and only differ in their risk aversion. The alternative polar case is that investors have the same utility functions but different probability beliefs. With identical risk preferences, all investors would transform a given p-density into the same RND and the pricing kernels from the three markets would be the same. But we have just seen that there are significant differences in the RNDs and pricing kernels in these markets.

If we assume investors all apply the same utility-based risk-neutralization, different pricing kernels must arise from different probability beliefs. SPY investors transform their estimate of the p-density, which we have assumed is lognormal with annualized mean equal to the riskless rate plus 5% and volatility calculated as shown in (A18) into $q_{SPY}(x)$, which produces the kernel $k_{SPY}(x)$. By inverting this transformation and applying it to the other two kernels, we can recover implied p-densities, $p_{SSO\text{-implied}}(x)$ and $p_{SDS\text{-implied}}(x)$. For example, if SPY investors believed the true p-density was $p_{SDS\text{-implied}}(x)$, they would risk-neutralize it to generate the same pricing kernel as we extracted from the SDS market.

Equation (6) shows this transformation.

$$p_{SDS\text{-}implied}(x) \;=\; \frac{e^{-r_f T} q_{SDS}(x)}{k_{SPY}(x)} \;=\; \frac{q_{SDS}(x)}{q_{SPY}(x)}\, p\text{-}density(x) \qquad (6)$$

Figure 5 displays the assumed p-density and the two implied p-densities derived by applying this formula to the data from May 2, 2008.

[insert Figure 5 about here]

Intuitively, one expects that the primary differences among SSO, SPY and SDS investors is in their expectations about the future level of the S&P index. Double long investors are more bullish than average, SPY investors have expectations similar to the market average, and SDS investors expect negative returns. Of course they may well differ in risk attitudes also, but there are many alternative ways to hedge market exposure, while options on these ETFs offer leveraged bets on specific portions of the returns distribution, which are not easily replicated in other ways.

This reasoning leads to hypotheses about how the implied-p densities will relate to one another. In particular, the p_{SSO}-implied density should show higher expected probabilities for large up moves and lower probabilities for losses than the other two, while the p_{SDS}-implied density should indicate greater chance of losses and less for profits on the index than the SSO or (our assumed) SPY p-densities. Although the hypothesized relationships are the same for moderate and large losses, for consistency with earlier tables, we continue to break the loss region into two parts.

In the region of moderate gains (0.0 to +1.5 $\hat{\sigma}$),

H4a: $p_{SSO\text{-implied}}(x) > p_{SPY}(x)$

H4b: $p_{\text{SDS-implied}}(x) < p_{\text{SPY}}(x)$
H4c: $p_{\text{SSO-implied}}(x) > p_{\text{SDS-implied}}(x)$

In the region of moderate losses (-1.5 to 0.0 $\hat{\sigma}$),
H4d: $p_{\text{SSO-implied}}(x) < p_{\text{SPY}}(x)$
H4e: $p_{\text{SDS-implied}}(x) > p_{\text{SPY}}(x)$
H4f: $p_{\text{SSO-implied}}(x) < p_{\text{SDS-implied}}(x)$

In the region of large losses (below -1.5 $\hat{\sigma}$),
H4g: $p_{\text{SSO-implied}}(x) < p_{\text{SPY}}(x)$
H4h: $p_{\text{SDS-implied}}(x) > p_{\text{SPY}}(x)$
H4i: $p_{\text{SSO-implied}}(x) < p_{\text{SDS-implied}}(x)$

Table 5 presents statistics on the average differences in the implied probabilities.

[Table 5 about here]

Under the assumption that holders of the different ETFs only differ in their returns expectations, SDS investors should anticipate a lower probability of a moderate rise in the S&P 500 and higher probability of a drop. These hypotheses, H4b, H4c, H4e and H4f, are all strongly confirmed in Table 5. However, the relationship between the p-density (assumed as) reflecting the expectations of SPY investors and the p-density implied from SSO options violates all three hypotheses, H2a, H2d, and H2g. The supposedly very bullish SSO investors are less eager than SPY investors to buy exposure on the upside and will pay more for payoffs that protect on the downside. This seemingly anomalous result could imply that many investors choose SSO over SPY because of the greater leverage, not because they are more bullish.

We expect that excluding risk preferences from consideration, double short investors should anticipate higher probability of a sharp market drop than investors who have chosen long market exposure. Table 5 indicates that this reasoning is correct on average for SDS versus SPY, although H4e is actually satisfied on only half of the days. But in this region, the implied p-density for double long SSO holders is significantly above that for SDS holders on average and on 57% of the sample days.

The results in this section are largely consistent with the expectations differences one expects from different sets of investors who choose long versus very long versus double short exposure to the market index, but neither of the polar cases fully captures the effects of their heterogeneity.

4. Response of Risk Neutral Densities and Implied Probabilities to Market Returns

So far we have explored differences in the risk neutral densities and associated curves from the three ETFs, averaged across the days in our sample. Many interesting questions relate to how they respond to realized returns in the market. Several different effects on both expectations and risk preferences might be at work. First, if the most recent return is positive, it tends to confirm the beliefs of bullish investors but is opposite to what the bears were expecting. Incorporating the new data point into their forecasting might lead all investors to modify their expectations about future returns and also their confidence in those expectations. In addition, the bulls have just become wealthier and the bears are poorer, which might well affect their risk appetites. Our data allows us to examine the differential impact that the same return shock to the underlying asset has on the risk neutral densities from ETFs with different exposures. Moreover, because we are extracting a density from the market and not just a single price or return, we can investigate effects on tail behavior, expected variance (i.e., uncertainty), and other properties of the RND along with the mean.

We cannot obtain the full q-density from market option prices, so the only way to compute its moments like mean and variance would be to append tails to the middle portion of the density in some way. As mentioned above, there are a variety of approaches in the literature to do this, but in all cases the tail shape is essentially imposed by the researcher's assumptions. In this paper, we wish to avoid such subjectivity and rely as much as possible on what can be extracted from observed market data, for example focusing on medians rather than means in Table 2. In this section, we will first report the impact of yesterday's return on the medians (as nonparametric proxies for means), the 25th and 75th percentiles (suggestive of tail behavior), and the interquartile spread (to reflect uncertainty over the expiration day level of the underlying index). We then explore the changes in the implied p-density in response to yesterday's market return. Unlike equation (6), which is based on an assumption that investors in each ETF risk-neutralize a given p-density in the same way, here we only assume that they each do this transformation the same way as they did it the day before.

To understand the dynamic behavior of the risk neutral densities we have calculated, one must keep in mind that the density is defined in terms of the standardized return from date t to date t+T, with 0 representing a (dividend-inclusive) return equal to the current riskless interest rate. In the canonical Black-Scholes asset market, where investors believe prices follow a semimartingale and they treat the equity risk premium and volatility as known constants, expectations about future returns or risk will not be affected by realized returns. The RND expressed in terms of standardized returns should show no change in the median or any quantile regardless of the observed return. But for investors with extrapolative expectations, a strong positive return predicts further excess returns and all RND quantiles will rise. On the other hand, if an investor's expectation is anchored on the underlying asset price rising to a particular level by option expiration, for example, a large positive return that takes it toward that level will reduce the rate of return expected in subsequent periods. In that case, after an unexpectedly large positive return, an SSO investor's RND may show a reduction in the median and the 75th percentile. With different reasoning but leading to the same result, the same positive return could make an SDS investor feel that the market was now even more overpriced than it was before and lead him to lower his expectations for the median return and left tail going forward.

A second kind of expectations response one might anticipate from real world investors is an adjustment in the degree of confidence they have in their prior beliefs, depending on whether the most recent return is consistent with or contrary to what they were expecting. A large positive return may make SSO investors more confident that they are right, causing a narrowing of the interquartile spread, while the spread for SDS investors would widen. If investors who choose leveraged exposures do so largely for speculative motives, these dynamic RND effects may be expected to be larger for SSO and SDS than for SPY.

In addition to these possible effects on their expectations, yesterday's return impacts the wealth of investors in these ETFs in different ways, which should alter their risk appetites. Other things equal, losing money should cause investors to want to reduce their exposure to the risk of further losses while a profit might encourage them to take on more risk. In that case, SSO and SPY holders with long exposure will become somewhat more aggressive on the upside following a strong positive return, while SDS investors will prefer to reduce risk exposure. This will show up as a relative increase in the RND quantiles on the right side (e.g., at the 75th percentile) for all three classes of investors. SSO and SPY holders want to increase their positive exposure and SDS holders want to reduce their negative exposure. However, we expect wealth effects to be of second order relative to the changes in expectations, since for most investors the actual daily returns in this market should be small relative to their total wealth.

This line of reasoning suggests a set of hypotheses for which the null is that the median and both the 25th and 75th percentiles of the q-density will be unaffected by the most recent return. Table 6 presents empirical results on the impact of the date t realized return on the risk neutral density for future returns. But given the large number of comparisons we will simply describe the results rather than formulating and testing specific hypotheses.

In presenting the results in this section, we introduce the following simple but potentially confusing notation. The realized return on some specific date t, calculated as $\log(S_t/S_{t-1})$ will be denoted r_t; r with no subscript is the average return over the remaining period up to option expiration, which also defines a state of the world covered by the RND on the expiration date. Both of these are expressed in terms of standard deviations relative to the expected value of the return. Thus, for example, $q_{SDS,t}(r)$ is the risk neutral density from the SDS ETF observed on date t for the event that the realized mean return from t to t+T is r standard deviations above the riskless rate. r_t is standardized relative to yesterday's volatility forecast, $\hat{\sigma}_{t-1}$, while r from t to t+T is standardized using the updated $\hat{\sigma}_t$.

[Table 6 about here]

The space of realized returns is split into six intervals, from less than -2.0 to more than +2.0 standard deviations. The first section of the table shows how the most recent observed return changes the median for each of the three ETF RNDs. If SSO investors have extrapolative expectations, a strong positive return on the S&P index will cause them to anticipate higher mean returns in the future and the quantiles of their RND should shift upward. The increase in wealth from today's profits should also raise the demand for upside exposure. These effects should occur with the SPY RND too, but to a lesser extent since their expectations are presumably less bullish than for SSO and the profit per dollar invested in their ETF is only half as large. By the same token, if yesterday's return was negative, the bulls may become less optimistic about sharp rises in the market in the immediate future and less eager to bear risk. SDS holders will have the opposite reactions: less bearish expectations and greater risk aversion leading to a rise in the median expected future return following a positive return, and more bearish expectations on a day after the market drops.

On the other hand, investors may not extrapolate a high positive return today into a higher mean for returns in the future. If SSO investors are Bayesians in updating their expectations for the expiration day level of the index, their date t+T forecasts will not rise by the full amount of today's increase, which will translate into lower expected mean returns than they predicted the day before and a negative change in the median. Extending this line of reasoning to SDS holders, it becomes the following: If the most recent return was in the direction they were expecting, and especially if it was large, they feel the market has probably overshot a little, so they reduce the rate of excess return they anticipate over the remaining period to expiration. In that case, the expectations effect and the wealth effect will operate in different directions.

The first portion of Table 6 shows that SSO investors appear to exhibit the first behavior with a high degree of statistical significance: a positive return increases the median of their RND and a negative return lowers it. SPY investors behave somewhat similarly: positive responses to large positive excess returns and a negative response to a negative return between -1.0 and -2.0 standard deviations, but the effects are smaller and less significant than for SSO. By contrast, SDS investors do exactly the opposite. Their RND median increases when today's return is negative and drops when it is positive.

The next two sections of Table 6 show how the 25th and 75th RND percentiles move in response to today's realized return. For SSO, the 25th percentile behaves the same way as the median, rising after a positive return and falling after a negative one. The situation for the 25th SPY percentile is almost the reverse of how the median behaved: negative (positive) returns appear to raise (lower) it, although only three of six buckets show significant effects at the 5 percent level in a one-tailed test. SDS follows the same pattern at the 25th and also at the 75th percentile as it did with the median: the RND moves in the opposite direction to the realized return. For both SSO and SPY, the 75th percentile is ambiguous, with only three significant coefficients, of which two were in the most positive return bucket that only contains 26 observations.

The last section of Table 6 is perhaps the most interesting, since it shows how the different ETF holders' confidence in their expectations about the future responds to current realized returns.

If an investor is wrong about the market's direction, his confidence in his prior belief should drop, and the interquartile spread between the 25th and the 75th percentile should widen. The larger the surprise, the bigger the effect should be. Similarly, if the most recent return confirms his directional belief, an option holder is likely to become more confident and narrow his range of uncertainty for future returns. So we expect negative average changes in the interquartile spreads for SSO and SPY following positive returns and for SDS holders following negative returns. Spreads should widen for SSO and SPY when the market falls and for SDS when it rises. Table 6 shows this pattern very strongly for the two leveraged ETFs, with SPY being largely consistent with SSO but not entirely so. Moreover, the size of the effect does increase with the size of the return, except for returns in the extreme tails.

These results still are subject to the problem that they combine both changes in expectations and in risk preferences. In the previous section, we computed pricing kernels and implied p-densities under the assumption that investors in each of the three ETFs either had the same estimates of the p-density as holders of the other ETFs or they had the same utility functions. Neither of these assumptions is particularly appealing. Much more plausible would be to assume that investors in the different ETFs have different risk preferences and different expectations, but each investor's date t relative utility of a future payoff under each return scenario is the same as on date t-1. In other words, their RNDs change from one day to the next because their estimates of the p-density change, but they risk-neutralize their updated p-densities in the same way as they did the day before.

Extending this reasoning over the whole sample would lead to constant risk preferences, which would allow p-densities to be estimated directly from the RNDs, but it is much stronger than we would like. Rather, we assume the investor's utility for a date t+T payoff in a state of the world corresponding to an average return of r % is only slowly time-varying and is largely independent of the predicted probability.

Specifically, adapting equation (2) to the standardized returns we have constructed, taking the SDS RND as an example, gives the following relationship on date t.

$$q_{SDS,t}(r) = p_{SDS,t}(r)\, k_{SDS,t}(r) \tag{7}$$

For a single ETF, the ratio of the RND for date t relative to date t-1 produces another kind of relative demand intensity similar to those we considered above.

$$RDI_{SDS,t,t-1}(r) = \frac{q_{SDS,t}(r)}{q_{SDS,t-1}(r)} = \frac{p_{SDS,t}(r)}{p_{SDS,t-1}(r)} \frac{k_{SDS,t}(r)}{k_{SDS,t-1}(r)}$$

Assuming SDS holders will risk-neutralize today's expected empirical density approximately the same way as they did yesterday, the ratio of pricing kernels in this expression will be 1, giving

$$RDI_{SDS,t,t-1}(r) \cong \frac{p_{SDS,t}(r)}{p_{SDS,t-1}(r)} \tag{8}$$

The ratio of the probability estimated on date t of a future standardized return of r, relative to that estimated probability on date t-1 is just the RDI constructed from that ETF's RNDs on the two dates. Table 7 reports those results for the three ETFs.

[Table 7 about here]

Table 7 uses the same breakdown of realized returns as in Table 6 and the same ranges of the future returns state space as in Tables 3-5. If there is no change in the p-density from t-1 to t, $RDI_{t,t-1}(r)$ will be 1.0 at every value of r. An $RDI_{t,t-1}(r)$ greater (less) than 1.0 indicates an increase (decrease) in the expected probability for a return to expiration following a realized r_t in the range for that column. Since both p-densities must integrate to 1.0, an increase in expected probability in one region must be offset by a decrease in another. (Note that this compensating change can be in the remote tails that we do not observe.) For each combination of date t realized return and future mean return, we report the average value of the $RDI_{t,t-1}$, the t-statistic on the hypothesis that it is equal to 1.0, and the fraction of days on which it was greater than 1.0.

Table 7 is clearly related to Table 6 and there are interesting comparisons to be made, but it is not easy to compare them directly. Table 6 shows how the 25th, 50th, and 75th quantiles of an RND change after a realized return of a given size, while Table 7 reports on how the expected probability that the future mean return will be in a specified range changes. Even so, the different perspective sheds additional light on some of the results we have just seen.

First consider the effect of a negative realized return on date t. For SPY investors, the expected probability of a large loss in the future goes up and so does the probability of a gain, while the chance of a subsequent moderate loss is reduced. SSO investors view any loss as an indication of higher probabilities of future losses and lower chance of future gains. And interestingly, SDS investors' predicted probabilities change in the same way as for SSO even though their RND quantiles responded in opposite directions in Table 6.

In the region of positive realized returns, large profits reduce the chance of large losses in the future for SPY investors but increase the chance of moderate losses. The effect a profit today on the probability of moderate gains in the future is ambiguous. For SSO, any gain today lowers the probability of losses in the future and increases the expected chance of a moderate gain. Again, surprisingly, SDS investors' expectations change the same way as for SSO.

5. Conclusion

Extracting an implied volatility from an option's market price has long been standard practice, although the result contains both the market's objective prediction of future volatility and also a volatility risk premium. A set of options with the same expiration and different strikes provides a major increase in market completeness and also allows (nearly) an entire probability density for the underlying asset to be extracted from market option prices. But like implied volatility, it is a risk-neutral density that impounds both the investors' true probability beliefs and their risk preferences.

Investors are observably quite heterogeneous, but most finance theory invokes the assumption of a representative investor with the property that a homogeneous population of such investors would produce exactly the same market prices as are seen in the real world. Applying this assumption to the stock index options market (in which truly homogeneous investors would not participate at all) consistently leads to the "puzzle" that the empirical pricing kernel is ill-behaved. The literature contains many efforts to explain its non-monotonic shape or to straighten it out, but one of the easiest and most plausible suggestions is that it can arise naturally from the combination of heterogeneous investors and limits to arbitrage. In this study, we have explored this heterogeneity by comparing and contrasting RNDs implied by options on ETFs with different exposures to the same underlying market index.

Analysis of RNDs from the three ETFs revealed highly significant differences among them, largely in the directions we anticipated, such as higher median expected returns for both the double long SSO and SPY relative to the double short SDS. But we still faced the problem that separating expectations and risk preferences unambiguously is not possible without additional assumptions. Even so, it is useful to compare the relative demand for payoffs under different expiration date contingencies by looking at the ratio of RNDs from the different ETF investor populations over different portions of the state space of returns, which we have called the Relative Demand Intensity. Relative to SPY investors, those who choose short exposure in the SDS ETF placed more value on payoffs in states with negative returns and less on positive returns, as expected, and the same held for SSO versus SDS for moderate sized returns. But for large negative index returns, aversion to the double-sized losses SSO investors would experience was apparently enough to offset belief that such bad returns were unlikely and RDI for SSO over SDS rose significantly above 1.0 for returns more than 1.5 standard deviations below the mean.

While such suggestive results can be obtained from studying RNDs, having three different densities over the same returns space let us investigate heterogeneous probability estimates versus risk preferences in three ways. We first considered the polar case in which all investors are assumed to have the same probability beliefs and only to differ in risk preferences, which allowed us to compare three different pricing kernels. The opposite polar case was to assume all investors had the same utility functions and would risk-neutralize any given empirical density in the same way, which allowed us to compare different implied p-densities for them.

Perhaps more interesting than the static comparisons between RNDs is the ability to look at how market returns change the way investors value payoffs under different future return scenarios. One great value of extracting an entire density function from market prices is that it becomes possible to look at how realized returns affect investors' confidence in their beliefs, as measured by the risk-neutral standard deviation, or the interquartile spread in this case.

Focusing on RND changes led to the third way we tried to look at true probability expectations separate from risk aversion. Instead of assuming all investors have the same utility functions, if each ETF's holder simply has the same risk preferences today as they did yesterday, their relative demand intensity for date t versus date t-1 reveals how their true probability beliefs have changed.

This first effort in this line of investigation has produced a number of interesting results and suggests that exploring investor heterogeneity through options on leveraged ETFs using the kinds of methodological tools we have proposed should be a very fruitful line for future research.

References

Ait-Sahalia Y. and A.W. Lo. (1998). "Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices." Journal of Finance 53, 499-547.

Ait-Sahalia Y. and A.W. Lo. (2000). "Nonparametric Risk Management and Implied Risk Aversion." Journal of Econometrics 94, 9-51.

Avellaneda, M. and S. Zhang (2009). "Path-dependence of Leveraged ETF returns." Working paper, New York University Courant Institute of Mathematical Sciences. Available at SSRN: http://dx.doi.org/10.2139/ssrn.1404708

Barone-Adesi, G., L. Mancini, and H. Shefrin (2013). "A Tale of Two Investors: Estimating Optimism and Overconfidence." Working paper, Santa Clara University. Available at SSRN: http://dx.doi.org/10.2139/ssrn.2319260 .

Bates, D. (2000). "Post-’87 Crash Fears in the S&P 500 Futures Option Market." Journal of Econometrics 94, 181-238.

Birru, J. and S. Figlewski (2012). "Anatomy of a Meltdown: The Risk Neutral Density for the S&P 500 in the Fall of 2008." Journal of Financial Markets 15, 151-180.

Bliss, R., Panigirtzoglou, N., 2004. Option-implied risk aversion estimates. Journal of Finance 59, 407–446.

Brown, David and Jens Jackwerth (2012). "The Pricing Kernel Puzzle: Reconciling Index Option Data and Economic Theory." in Derivative Securities and Modeling, J. Batten and N. Wagner, eds., Emerald, UK.

Breeden, D. and R. Litzenberger (1978). "Prices of state-contingent claims implicit in option prices." Journal of Business 51, 621–652.

Breeden, D. and R. Litzenberger (2013). "Central Bank Policy Impacts on the Distribution of Future Interest Rates." Working paper, available at: http://www.dougbreeden.net/uploads/Breeden_Litzenberger_with_Postscript_Central_Bank_Policy_Impacts_9_20_2013.pdf

Eriksson, A., Ghysels, E., Wang, F., (2009). The normal inverse Gaussian distribution and the pricing of derivatives. Journal of Derivatives 16, 23–37.

Federal Reserve Bank of Minneapolis, (2014). "Estimates of the Future Behavior of Asset Prices." available at: https://www.minneapolisfed.org/banking/rnpd/.

Fernandez, P. and J. del Campo (2010). "Market Risk Premium used in 2010 by Analysts and Companies: a survey with 2,400 answers." Working paper, IESE Business School.

Figlewski, S. (2010). "Estimating the implied risk neutral density for the U.S. market portfolio." In: Bollerslev, T., Russell, J.R., Watson, M. (Eds.),Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle. Oxford University Press, Oxford, UK.

Gemmill, G., Saflekos, A., 2000. How useful are implied distributions? Evidence from stock-index options. Journal of Derivatives 7, Spring, 83–98.

Harrison, J. M. and D.M. Kreps (1979). "Martingales and arbitrage in multiperiod securities markets." Journal of Economic Theory 20, 381-408.

Hens, T. and C. Reichlin (2013). "Three Solutions to the Pricing Kernel Puzzle." Review of Finance,17, 1029-64.

Jackwerth, J.C. (2004). Option-Implied Risk-Neutral Distributions and Risk Aversion. Charlotteville: Research Foundation of AIMR.

Jackwerth, J. and M. Rubinstein (1996). "Recovering Probability Distributions from Option Prices." Journal of Finance 51, 1611-31.

Liu, X., Shackleton, M.B., Taylor, S.J., Xu, X., 2007. Closed form transformations from risk-neutral to real-world distributions. Journal of Banking and Finance 31, 1501–1520.

Melick W.R. and C.P. Thomas (1997). "Recovering an Asset's Implied PDF from Option Prices: An Application to Crude Oil during the Gulf Crisis." Journal of Financial and Quantitative Analysis 32, 91-115.

Rompolis, L. and E. Tzavalis (2008). "Recovering Risk Neutral Densities from Option Prices: A New Approach." Journal of Financial and Quantitative Analysis 43, 1037-54.

Rosenberg, J. and R. Engle (2002). "Empirical Pricing Kernels." Journal of Financial Economics 64, 341-72.

Ross, S. (1976). "Options and Efficiency." Quarterly Journal of Economics 90, 75-89.

Ross, S. (2014). "The Recovery Theorem." Journal of Finance forthcoming, DOI: 10.1111/jofi. 12092.

Shefrin, H. (2008). A Behavioral Approach to Asset Pricing, 2nd ed. Academic Press: Burlington MA.

de Vincent-Humphreys, R. and J. Noss, (2012). "Estimating probability distributions of future assets prices: empirical transformations from option-implied to real-worl density functions." Bank of England Working Paper No. 455.

Appendix

This Appendix provides the details on extracting a risk neutral density from a set of options prices; the specific sample and data selection criteria we have used in this study; the procedure for converting risk neutral densities over the terminal values of the three ETFs into densities over the true underlier for all three ETFs, which is the level of the Standard and Poor's 500 stock index on option expiration day; and the procedure for converting each day's RND that depends on the current volatility, riskfree interest rate, and number of days to option expiration into a common base defined in terms of a number of standard deviations around the expected mean return on the index, for the p-density, and around the riskless rate for the q-density.

Extracting the RND

In the following, the symbols C, S, x, r_f, and T all have the standard meanings of option valuation: C = call price; S = time 0 price of the underlying asset; x = exercise price; r_f = riskless interest rate; T = option expiration date, which is also the time to expiration. We will also use q(x) = the risk neutral probability density function, also denoted RND, and Q(x) = the risk neutral distribution function.

The value of a call option is the expected value of its payoff on the expiration date T, discounted back to the present. Under risk neutrality, the expectation is taken with respect to the risk neutral probabilities and discounting is at the risk-free interest rate.

$$C = e^{-r_f T} \int_x^{\infty} (S_T - x) q(S_T) dS_T \tag{A1}$$

Taking the partial derivative in (A1) with respect to the strike price x and solving for the risk neutral distribution Q(x) yields:

$$Q(x) \quad = \quad e^{r_f T} \frac{\partial C}{\partial x} + 1 \quad . \tag{A2}$$

Taking the derivative with respect to x a second time gives the RND function:

$$q(x) \quad = \quad e^{r_f T} \frac{\partial^2 C}{\partial x^2} \tag{A3}$$

In practice, we approximate the solution to (A3) using finite differences. In the market, option prices for a given maturity T are available at discrete exercise prices that can be far apart. To generate smooth densities, we interpolate to obtain option values on a denser set of equally spaced strikes.

Let $\{x_1, x_2, ..., x_N\}$ represent the set of strike prices, ordered from lowest to highest, for which we have simultaneously observed option prices, and denote the price of a call option with strike price x_n as C_n. To estimate the probability in the left tail of the risk neutral distribution up to x_2, we approximate $\frac{\partial C}{\partial x}$ at x_2 and compute

$$Q(x_2) \cong e^{r_f T} \frac{C_3 - C_1}{x_3 - x_1} + 1 \tag{A4}$$

The probability in the right tail from x_{N-1} to infinity is approximated by,

$$1 - Q(x_{N-1}) \cong 1 - \left(e^{r_f T} \frac{C_N - C_{N-2}}{x_N - x_{N-2}} + 1 \right) = -e^{r_f T} \frac{C_N - C_{N-2}}{x_N - x_{N-2}} \quad . \tag{A5}$$

The approximate the density at a strike x_n, $q(x_n)$, is given by:

$$q(x_n) \cong e^{r_f T} \frac{C_{n+1} - 2C_n + C_{n-1}}{(\Delta x)^2} \tag{A6}$$

For the interpolation step, we use the mid-quote prices from the options market and transform them into equivalent Black-Scholes implied volatilities. This greatly reduces the impact of heteroskedasticity that would be a large problem if we simply tried to smooth market option prices that can differ by up to two orders of magnitude between in the money and out of the money contracts. We fit a 4^{th} degree polynomial through the volatility smile curve for each of the ETFs. Although ETF options are American, the Black-Scholes formula is only used here as a kind of transform to allow interpolation in volatility rather than strike price space. The inverse transform is then applied to the interpolated volatility smile to calculate a dense set of option prices from which the risk neutral probabilities are calculated using Equation (A6).

As is standard practice, we use only out of the money and at the money put and call options for our analysis. In the money options trade close to their intrinsic values and have higher bid/offer spreads, so the information they provide about probabilities is limited and noisy. From the put-call parity equation, it is easy to show that the second partial derivatives of the call and put value functions are equal at the same strike price, so (A3)-(A6) can be applied directly to puts.

However, it is common to find that despite the arbitrage constraint imposed by put-call parity, even at the money puts and calls can trade at somewhat different implied volatilities. To avoid an artificial jump in the RND when moving from OTM puts to OTM calls, we blend the implied volatilities from puts and calls over a ± 5 point range around the ATM forward level to produce a smooth transition.

Data Sample

SDS and SSO started trading in the middle of 2006, however, options on these leveraged ETFs only were introduced towards the end of 2007. For this study, we use the interval between October 2007 and March 2014, giving us 1615 potential observations for SPY, 1614 for SSO and 1611 for SDS. There were several periods of market disruption within this time span in which option quotes were not available for a wide enough range of strike prices or were too noisy to produce proper RNDs, typically because price quotes violated no-arbitrage conditions. The final sample contains 1440 usable RNDs for each of the three ETFs.

It is important to use price quotes rather than transactions prices in extracting risk neutral densities because simultaneously observed prices are required across the whole range of strikes, which is never possible with just trades. We use end-of-day market bid and ask quotes for SPY, SDS and SSO from OptionMetrics. The bid price was required to be greater than 0.05. The maturity date is always in the next month, which results in a time to maturity between 16 and 52 days. To calculate the forward levels for the ETFs, dividend and riskless interest rate data are obtained from OptionMetrics.

Rescaling Densities

Each of the three risk neutral densities has a different domain, determined by the market prices of the three ETFs. For instance, the following chart shows the RNDs for SDS, SSO and SPY as functions of the strike prices for options traded on May 2, 2008.

Figure A1: RNDs from SSO, SPY, and SDS before rescaling

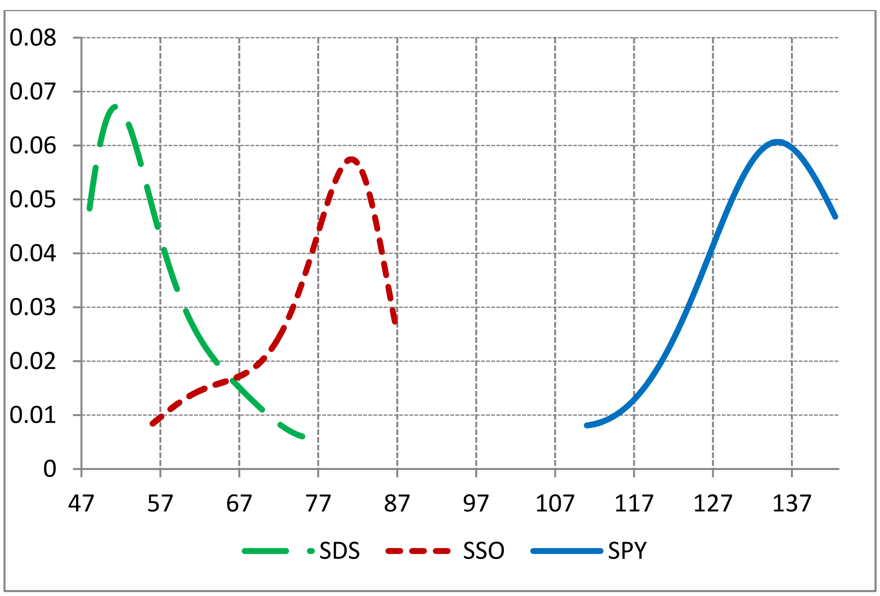

Converting to a common domain based on the S&P 500 index, which is the true underlying asset for all three ETFs, is not a trivial task because of the need for a convexity correction. Each ETF attempts to match the realized return on the S&P index times the ETF's leverage factor each day. The problem is not unlike trying to delta-hedge an option using daily rebalancing. Hedging error, or tracking error in this case, is a function of the curvature of the value function and the realized volatility.

An expression for the expected log return on a leveraged ETF as a function of the return on the underlying index can be easily derived for the case of a lognormal diffusion. We will apply that approximation to convert the RNDs based on ETF prices to a common base in terms of S&P 500 index returns.

Let S be the underlying index, with instantaneous mean r and volatility σ.

$$dS \;=\; rS\,dt \;+\; \sigma S\,dz \qquad \text{(A7)}$$

d logS is given by Ito's Lemma:

$$d\ logS \;=\; \frac{\partial\ logS}{\partial\ S}\,dS \;+\; \frac{1}{2}\,\frac{\partial^2\ logS}{\partial\ S^2}\,(dS)^2$$

With

$$\frac{\partial\ logS}{\partial\ S} = \frac{1}{S}\ ; \qquad \frac{\partial^2\ logS}{\partial\ S^2} \;=\; -\,\frac{1}{S^2}\ ; \qquad (dS)^2 \;=\; \sigma^2 S^2\,dt$$

Which yields

$$d\ logS \;=\; \left(r - \frac{\sigma^2}{2}\right)dt \;+\; \sigma\,dz \qquad \text{(A8)}$$

Over a horizon of length T, we have

$$\log\left({S_T}/{S_0}\right) \;=\; \int_0^T d\log S \;=\; \left(r - \frac{\sigma^2}{2}\right)T \;+\; \sigma\sqrt{T}\,Z \;=\; \log(\,1 + \,R_S)$$

Where Z is a N(0,1) random variable R_S is the realized return on the index and R_E is the realized return on the ETF.

Taking the exponential gives the gross return on the index

$$1 + \,R_S \;=\; e^{\left(r - \frac{\sigma^2}{2}\right)T \,+\, \sigma\sqrt{T}\,Z} \qquad \text{(A9)}$$

We are interested in an ETF that is set up to return L times the instantaneous return on the index at all points in time. Let E be the price of the ETF.

$$\frac{d\,E}{E} \;=\; L\,\frac{dS}{S} = \;L\,r\,dt \;+\; L\sigma dz$$

$$dE \;=\; EL\,r\,dt \;+\; EL\sigma dz \qquad \text{(A10)}$$

Applying Ito's Lemma as before,

$$d\;logE \;=\; \left(L\,r - \frac{L^2\sigma^2}{2}\right)dt \;+\; L\,\sigma\,dz$$

Integrating gives

$$\log\left({E_T}/{E_0}\right) \;=\; \int_0^T d\log E \;=\; L\,\left(r - \frac{L\sigma^2}{2}\right)T \;+\; L\,\sigma\,\sqrt{T}\,Z$$

$$=\; L\left(\left(r - \frac{\sigma^2}{2}\right)T \;+\; \sigma\sqrt{T}\,Z\right) + \frac{L\sigma^2}{2}T \;-\; \frac{L^2\sigma^2}{2}T$$

$$=\; log\left({S_T}/{S_0}\right)L \;+\; (L - L^2)\frac{\sigma^2}{2}T \qquad \text{(A11)}$$

Taking the exponential gives the gross return on the ETF,

$$1 \;+\; R_E \;=\; (1 \;+\; R_S)^{\,L}\,e^{(L-L^2)\frac{\sigma^2}{2}T} \qquad \text{(A12)}$$

From this equation we can see that the leveraged ETF has the highest return when the volatility of the underlying is zero. Any other level of volatility would give a lower return than the term

$(1 \;+\; R_S)^{\,L}$.

The SSO "Double Long" ETF has L = 2 and the SDS "Double Short" ETF has L = -2. We want to find the RND for future SPY levels that is implied from the RND extracted from the levered ETF option prices.

Let E_T be the future level of the ETF at date T and E_0 be its initial level. The RND for E_T is extracted from options on the leveraged ETF. We want to use the RND for E_T to compute an RND for S_T that is implied in the ETF options.

We use the formula for a transformation of the density for a variable x, $f_x(x)$, to the density $f_Y(Y)$, for Y, which is a function of x: Y = g(x):

$$f_Y(Y) = \left|\frac{d}{dY}g^{-1}(Y)\right| f_X(g^{-1}(Y))$$

The ETF RND relates ETF option prices to the density for the future level of the ETF. Consider first the transformation from the empirical density for the SPx index to the empirical density for the levered ETF.

From (A12) we have

$$E_T \;=\; E_0\, ({}^{S_T}/_{S_0})^{\,L}\; e^{(L-L^2)\frac{\sigma^2}{2}T} \;=g(S_T) \tag{A13}$$

Solving for $g^{-1}(E_T)$,

$$S_T \;=\; S_0\, ({}^{E_T}/_{E_0})^{\frac{1}{L}}\, e^{(L-1)\frac{\sigma^2}{2}T} \;=\; g^{-1}(E_T) \tag{A14}$$

Taking the derivatives in (A14) gives:

$$\frac{d}{dE_T}\; g^{-1}(E_T) \;=\; \frac{S_0}{L\, E_0^{1/L}}\; e^{(L-1)\frac{\sigma^2}{2}T}\; E_T^{(\frac{1}{L}-1)} \tag{A15}$$

If the empirical density of the underlying index $f_x(S_T)$ is lognormal with mean rT and volatility $\sigma\sqrt{T}$, the empirical density of the ETF with leverage factor L, $f_L(E_T^L)$ at some value E^{L*} is given by

$$f_L(E_T^{L*}) \left| \frac{S_0}{L\, E_0^{1/L}}\; e^{(L-1)\frac{\sigma^2}{2}T}\; E_T^{L*(\frac{1}{L}-1)} \right|^{-1} \;=f_X(S_T^*) \tag{A16}$$

where $S_T^* = g^{-1}(E_T^{L*})$ is the index level corresponding to a date T ETF value E^{L*}.

For example, applying the transformation in (A16), to the RNDs for May 2, 2008 gives the following plot, on the common base of the S&P 500 index divided by 10, which is the price of the SPY ETF:

Figure A2: RNDs from SSO, SPY, and SDS after converting to a common price basis

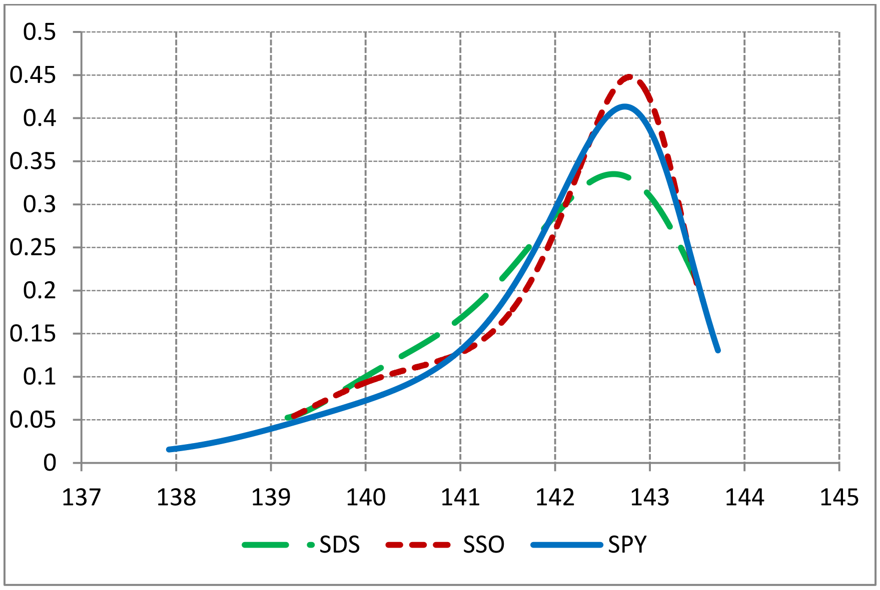

<u>Normalizing the Domain</u>

With the normalization just described it is possible to compare the RNDs for the three ETFs on a given day, but not across days. The RNDs on some date t are extracted from options that mature on the next expiration date, whose prices will depend on the number of days to expiration, the date t market interest rate, and investors' volatility forecasts. We need to transform them again to allow comparisons across the full sample.

We do this with a second transformation to a common base expressed in terms of return standard deviations. Our assumption is that under the p-density for every date the log return is normal, with mean ($r_{f,t}$ + .05) T/365 and standard deviation $\widehat{\sigma_t}\sqrt{T/365}$ where $r_{f,t}$ is the interpolated annual riskless rate on date t, T is the number of calendar days to expiration, and $\widehat{\sigma_t}$ is the market's volatility forecast.

To obtain a value for $\widehat{\sigma_t}$ we assume that in forecasting the volatility over the time to option maturity, investors take into account both historical volatility from recent returns as well as implied volatility represented by the VIX index. More precisely, for each date t for which we extract a q-density from options maturing T days later, (i.e., on date t+T), we run the following regression on past data.

$$v(\tau,T) = c + a\,V_{63\,day}(\tau) + b\,VIX(\tau) + \varepsilon(\tau,T) \qquad \text{(A17)}$$

where v(τ,T) is the volatility realized over the subsequent T days, from date τ+1 to τ+T;

$V_{63\ day}$ (τ) is realized volatility over the previous 63 trading days; and VIX(τ) is the VIX index on date τ. In each case, the regression sample period begins on Jan. 2, 1990 and ends on date t-T. A forecast of (annualized) volatility for the period t+1 to t+T is then computed from the estimated regression coefficients,

$$\hat{\sigma}_t = \hat{c} + \hat{a}\, V_{63\ day}(t) + \hat{b}\, VIX(t) \qquad \text{(A18)}$$

Hatted right hand side variables are the coefficient estimates from the (A17) regression fitted on all past data up to date t. A given curve is approximated by a histogram of discrete S_{t+T} values and associated probabilities. Each S_{t+T} is converted into a standardized value expressed as a number of standard deviations:

$$S_{t+T} \;\rightarrow\; \frac{\log(S_{t+T}/S_t) - \left(r_{f,t} + 0.05\right) T/365}{\widehat{\sigma_t}\sqrt{T/365}}$$

(A19)

The new common x-axis is divided into buckets of width 0.0025, with the original probabilities appropriately distributed into the new buckets. The axis runs from -4.0000 to +4.0000 in steps of 0.0025, a total of 3200 buckets, which is wider than any of the actual curves in the sample.

Table 1: Valid RND Observations in Different Ranges of Standardized Return

The data sample consists of call and put options on three S&P 500-based ETFs, with maturity between 16 and 52 days, from 10/23/07 to 3/21/14. A volatility forecast combining historical volatility and the current VIX index is constructed for each day using eq. (A18). Each available exercise price is converted into the equivalent number of standard deviations relative to an annualized return equal to the riskless interest rate over the period to expiration. Some days were eliminated due to insufficient data, as described in the Appendix. The table shows the number of sample days that produce valid RNDs over the standardized range.

Standard deviations	SPY	SSO	SDS
-3	1434	1022	927
-2.5	1440	1127	1188
-2	1440	1261	1348
-1.5	1440	1390	1428
-1	1440	1422	1437
-0.5	1440	1435	1438
0	1440	1440	1439
0.5	1440	1440	1389
1	1440	1393	1021
1.5	1034	406	161
2	47	3	4
2.5	1	0	0
3	0	0	0

Table 2: Hypotheses on Medians of Risk Neutral Densities

Risk neutral densities are extracted from the options data and rebased to a common basis in terms of standard deviations of return as described in the Appendix. The top portion of the table shows the average and standard deviation of the medians of the RNDs across the 1440 sample days. The second portion reports statistics on the differences between the medians from different ETFs.

Levels	SPY	SSO	SDS
nobs	1440	1440	1440
average	0.216	0.191	-0.175
standard deviation	0.052	0.054	0.301
Differences: 1st minus 2nd	SSO-SPY	SSO-SDS	SPY-SDS
fraction positive	0.303	0.974	0.982
average difference	-0.025	0.366	0.391
standard deviation	0.042	0.312	0.308
t-statistic	-22.519	44.552	48.269

Table 3: Relative Demand Intensities in Regions of the Returns Distribution

The table reports statistics on the ratio of the RNDs for pairs of ETFs in different regions of the returns space. The ratio is averaged over the specified region for each day. The t-statistic tests the hypothesis that the mean of these averaged ratios is equal to 1.0.

	Relative Demand Intensity	SSO/SPY	
	-3.0 ≤ r < -1.5	**-1.5 ≤ r < .0**	**0 ≤ r < 1.5**
nobs	1390	1440	1440
fraction > 1.0	0.622	0.553	0.500
average	1.035	1.025	0.997
standard deviation	0.119	0.085	0.052
t-stat (diff from 1.0)	10.942	11.137	-1.948
	Relative Demand Intensity	**SDS/SPY**	
	-3.0 ≤ r < -1.5	**-1.5 ≤ r < .0**	**0 ≤ r < 1.5**
nobs	1428	1440	1439
fraction > 1.0	0.500	0.797	0.268
average	1.009	1.096	0.949
standard deviation	0.175	0.104	0.066
t-stat (diff from 1.0)	1.935	35.088	-29.491
	Relative Demand Intensity	**SSO/SDS**	
	-3.0 ≤ r < -1.5	**-1.5 ≤ r < .0**	**0 ≤ r < 1.5**
nobs	1383	1440	1439
fraction > 1.0	0.573	0.223	0.726
average	1.058	0.942	1.055
standard deviation	0.240	0.073	0.085
t-stat (diff from 1.0)	8.899	-29.994	24.544

Table 4: Pricing Kernel Differences in Regions of the Returns Distribution

Pricing kernels are computed for the three ETFs by dividing the RND at each return level by the same estimate of the p-density, constructed assuming lognormality, annualized mean return equal to the riskless rate plus 5% and volatility estimated using (A18). The table reports statistics on the differences between the pricing kernels over three regions of the returns space.

	SSO kernel - SPY kernel		
	-3.0 ≤ r < -1.5	**-1.5 ≤ r < .0**	**0 ≤ r < 1.5**
nobs	1440	1440	1440
fraction > 0	0.622	0.553	0.500
average	0.040	0.016	-0.007
standard deviation	0.261	0.059	0.070
t-stat	5.753	9.983	-3.637
	SDS kernel - SPY kernel		
	-3.0 ≤ r < -1.5	**-1.5 ≤ r < .0**	**0 ≤ r < 1.5**
nobs	1440	1440	1440
fraction > 0	0.500	0.797	0.268
average	-0.046	0.067	-0.069
standard deviation	0.416	0.072	0.085
t-stat	-4.171	35.296	-30.796
	SSO kernel - SDS kernel		
	-3.0 ≤ r < -1.5	**-1.5 ≤ r < .0**	**0 ≤ r < 1.5**
nobs	1440	1440	1440
fraction > 0	0.573	0.223	0.726
average	0.065	-0.051	0.059
standard deviation	0.372	0.062	0.074
t-stat	6.668	-30.866	29.969

Table 5: Differences in the Implied p-densities over

Regions of the Returns Distribution

Implied p-densities are computed for the SSO and SDS ETFs by dividing the RND at each return level by the same pricing kernel as constructed for SPY investors (the "actual" p-density derived using (A18) and (A19)). This assumes holders of the three ETFs all have the same risk preferences and differ only in their expected returns distributions. The table reports statistics on the differences between the p-densities over three regions of the returns space.

	SSO implied p-density - actual p-density		
	$-3.0 \leq r < -1.5$	**$-1.5 \leq r < .0$**	**$0 \leq r < 1.5$**
nobs	1440	1440	1440
fraction > 0	0.622	0.553	0.500
average	0.0021	0.0053	-0.0011
standard deviation	0.0071	0.0230	0.0176
t-stat	11.232	8.808	-2.387
	SDS implied p-density - actual p-density		
	$-3.0 \leq r < -1.5$	**$-1.5 \leq r < .0$**	**$0 \leq r < 1.5$**
nobs	1440	1440	1440
fraction > 0	0.500	0.797	0.268
average	0.0017	0.0254	-0.0164
standard deviation	0.0092	0.0280	0.0227
t-stat	6.882	34.420	-27.327
	SSO implied p-density - SDS implied p-density		
	$-3.0 \leq r < -1.5$	**$-1.5 \leq r < .0$**	**$0 \leq r < 1.5$**
nobs	1440	1440	1440
fraction > 0	0.573	0.223	0.726
average	0.0009	-0.0202	0.0144
standard deviation	0.0098	0.0254	0.0204
t-stat	3.469	-30.162	26.687

Table 6: Change in RND Quantiles and Interquartile Spread as a Function of Return

The RND for each ETF is computed for date t-1 and date t at each return level. The table reports averages across observation dates of the changes in the median, the 25th and 75th percentiles, and the interquartile spread. These are separated into buckets based on the date t-1 standardized return.

	$r_t < -2.0$	$-2.0 \le r_t < -1.0$	$-1.0 \le r_t < 0.0$	$0.0 \le r_t < 1.0$	$1.0 \le r_t < 2.0$	$2.0 \le r_t$
	ΔMedian					
q_{SPY}	0.0007	-0.0067	-0.0004	-0.0023	0.0117	0.0149
	(0.107)	(-2.013)	(-0.268)	(-1.656)	(3.424)	(2.063)
q_{SSO}	-0.0430	-0.0269	-0.0039	0.0039	0.0242	0.0500
	(-4.041)	(-6.321)	(-2.054)	(2.590)	(5.730)	(4.659)
q_{SDS}	0.2190	0.2397	0.0785	-0.0607	-0.1622	-0.1898
	(3.045)	(5.831)	(3.926)	(-4.244)	(-5.109)	(-1.587)
	Δ25th percentile					
q_{SPY}	0.0029	0.0001	0.0021	-0.0020	0.0024	-0.0087
	(0.814)	(0.077)	(1.990)	(-2.219)	(0.990)	(-2.177)
q_{SSO}	-0.0129	-0.0100	-0.0019	0.0012	0.0078	0.0154
	(-2.112)	(-3.832)	(-1.432)	(1.205)	(2.725)	(2.432)
q_{SDS}	0.1043	0.0814	0.0080	-0.0124	-0.0536	-0.0335
	(3.407)	(3.302)	(0.639)	(-1.376)	(-3.005)	(-0.367)
	Δ75th percentile					
q_{SPY}	-0.0058	-0.0006	0.0009	0.0014	-0.0015	-0.0168
	(-1.330)	(-0.317)	(0.911)	(1.587)	(-0.849)	(-2.895)
q_{SSO}	-0.0013	0.0008	0.0017	0.0007	-0.0055	-0.0199
	(-0.253)	(0.358)	(1.530)	(0.710)	(-2.612)	(-3.752)
q_{SDS}	0.1012	0.0563	0.0047	-0.0097	-0.0385	-0.0610
	(3.248)	(4.024)	(0.578)	(-1.708)	(-3.081)	(-1.050)
	ΔInterquartile spread					
q_{SPY}	-0.0066	0.0061	0.0013	0.0036	-0.0132	-0.0318
	(-0.713)	(1.305)	(0.591)	(1.924)	(-3.051)	(-2.862)
q_{SSO}	0.0416	0.0281	0.0058	-0.0037	-0.0297	-0.0699
	(3.078)	(4.803)	(2.273)	(-1.717)	(-5.407)	(-5.305)
q_{SDS}	-0.1568	-0.1918	-0.0658	0.0493	0.1283	0.1249
	(-2.630)	(-5.505)	(-3.995)	(4.099)	(5.021)	(1.274)
nobs	48	133	439	568	165	26

Table 7: Ratio of Implied p-density to Previous Day as a Function of Return

For each ETF, the ratio of the RND for date t-1 and date t is computed at each return level. Assuming investors in a given ETF risk-neutralize their density forecasts the same way on both days, this Relative Demand Intensity gives the ratio of expected returns probabilities. The ratios for each day are averaged across three ranges of the returns space corresponding to large losses, moderate losses, or moderate profits on the S&P index over the life of the options. The table reports average probability ratios across observation dates, separated into buckets based on the date t standardized return. t-statistics relate to the hypothesis that the ratio is 1.0.

		$r_t < -2.0$	$-2.0 \le r_t < -1.0$	$-1.0 \le r_t < 0.0$	$0.0 \le r_t < 1.0$	$1.0 \le r_t < 2.0$	$2.0 \le r_t$
		Large Loss: Standardized returns from -3.0 to -1.5					
SPY	average	1.0063	1.0168	1.0095	0.9990	0.9886	0.9606
	t-stat	0.594	3.115	3.560	-0.401	-2.023	-3.853
	fraction > 1.0	0.551	0.607	0.552	0.468	0.407	0.226
SSO	average	1.0779	1.0383	1.0131	1.0003	0.9780	0.9500
	t-stat	3.799	4.235	2.178	0.058	-3.166	-2.889
	fraction > 1.0	0.702	0.677	0.531	0.470	0.334	0.239
SDS	average	1.0480	1.0105	1.0411	1.0270	0.9791	0.9618
	t-stat	1.646	0.647	3.777	2.807	-1.871	-1.211
	fraction > 1.0	0.494	0.511	0.549	0.496	0.427	0.398
		Moderate Loss: Standardized returns from -1.5 to 0.0					
SPY	average	0.9866	0.9958	0.9963	1.0065	1.0020	1.0164
	t-stat	-2.200	-1.341	-2.091	3.976	0.453	1.892
	fraction > 1.0	0.410	0.426	0.457	0.571	0.596	0.675
SSO	average	1.0302	1.0129	1.0066	0.9996	0.9945	0.9751
	t-stat	1.586	2.348	2.011	-0.153	-0.933	-1.963
	fraction > 1.0	0.576	0.539	0.526	0.523	0.523	0.541
	NOBS						
SDS	average	1.0330	1.0248	1.0133	0.9964	0.9813	0.9672
	t-stat	3.200	4.996	4.795	-1.749	-4.544	-2.963
	fraction > 1.0	0.631	0.612	0.561	0.464	0.346	0.354
		Moderate Gain: Standardized returns from 0.0 to 1.5					
SPY	average	1.0078	1.0052	1.0034	0.9978	1.0027	0.9974
	t-stat	2.153	2.638	2.796	-2.275	1.015	-0.532
	fraction > 1.0	0.533	0.529	0.526	0.452	0.467	0.473
SSO	average	0.9987	0.9876	0.9967	1.0022	1.0188	1.0338
	t-stat	-0.177	-2.972	-1.685	1.225	4.452	3.865
	fraction > 1.0	0.437	0.437	0.467	0.481	0.563	0.630
SDS	average	1.0035	0.9964	0.9935	1.0045	1.0107	1.0266
	t-stat	0.447	-0.792	-2.606	2.300	2.551	1.819
	fraction > 1.0	0.526	0.489	0.450	0.513	0.535	0.715
	NOBS	48	133	439	568	165	26

Figure 1: Pricing Kernel for the S&P 500 Index Extracted from 50-day SPY Options, May 2, 2008

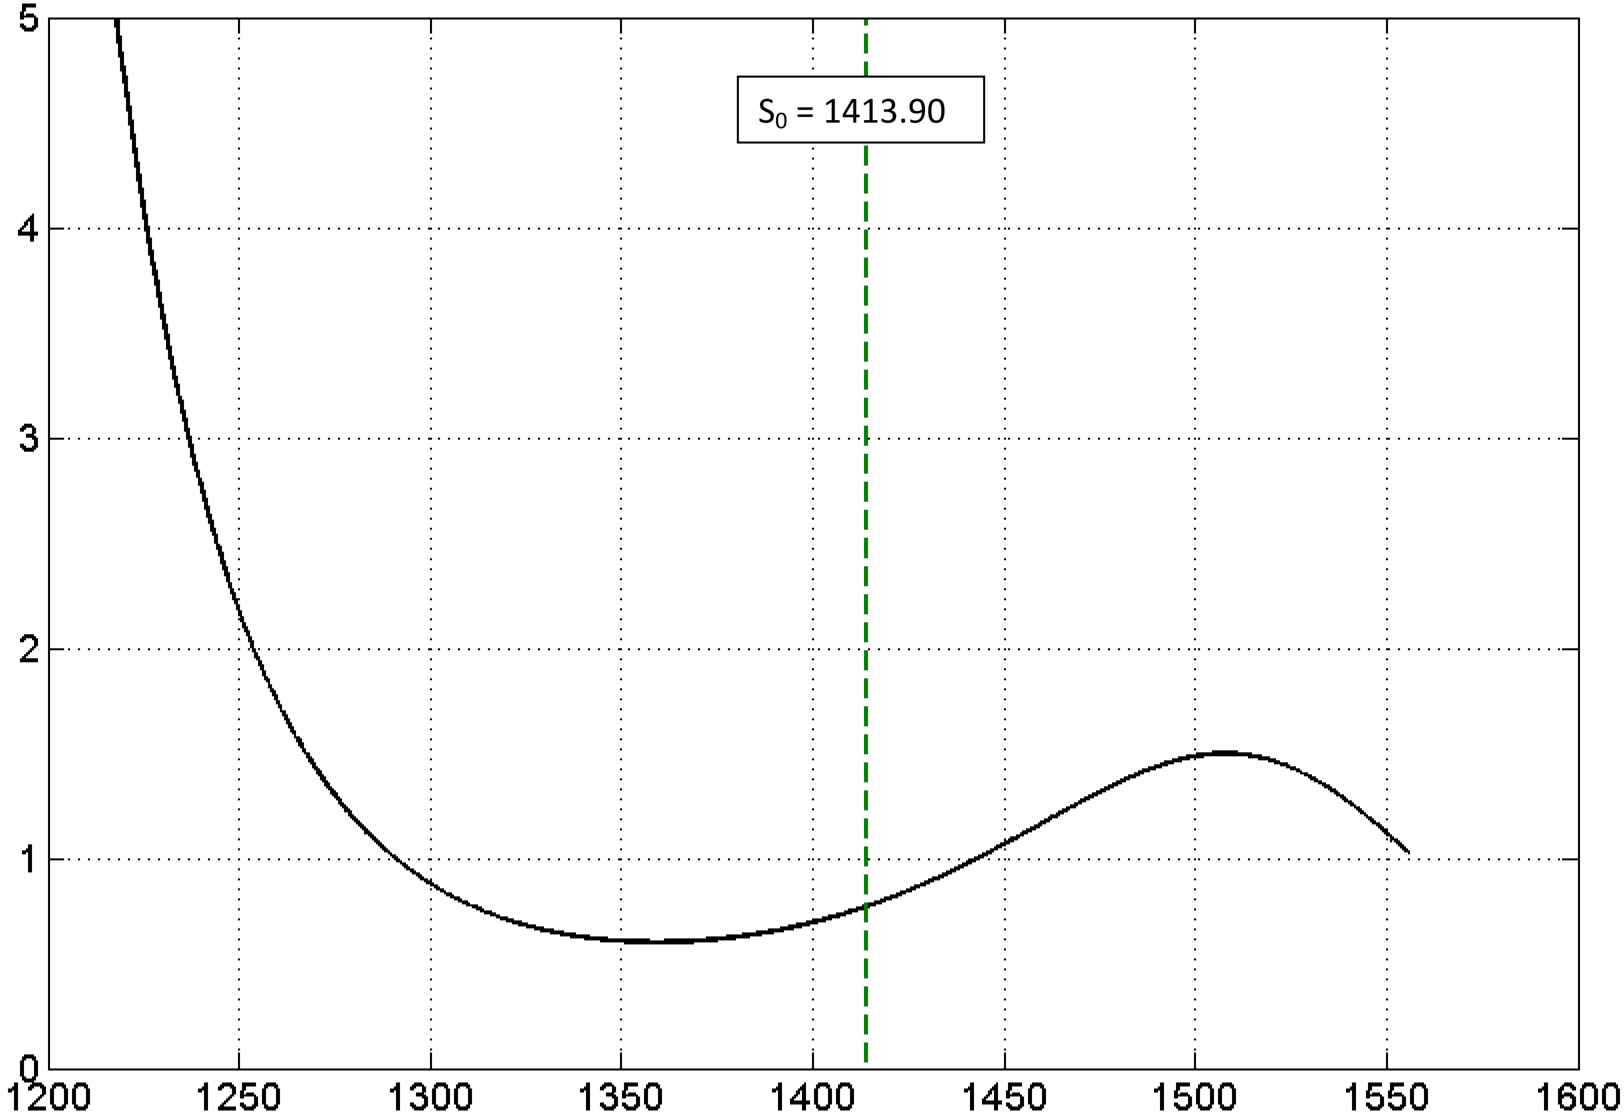

Figure 2: RNDs on May 2, 2008

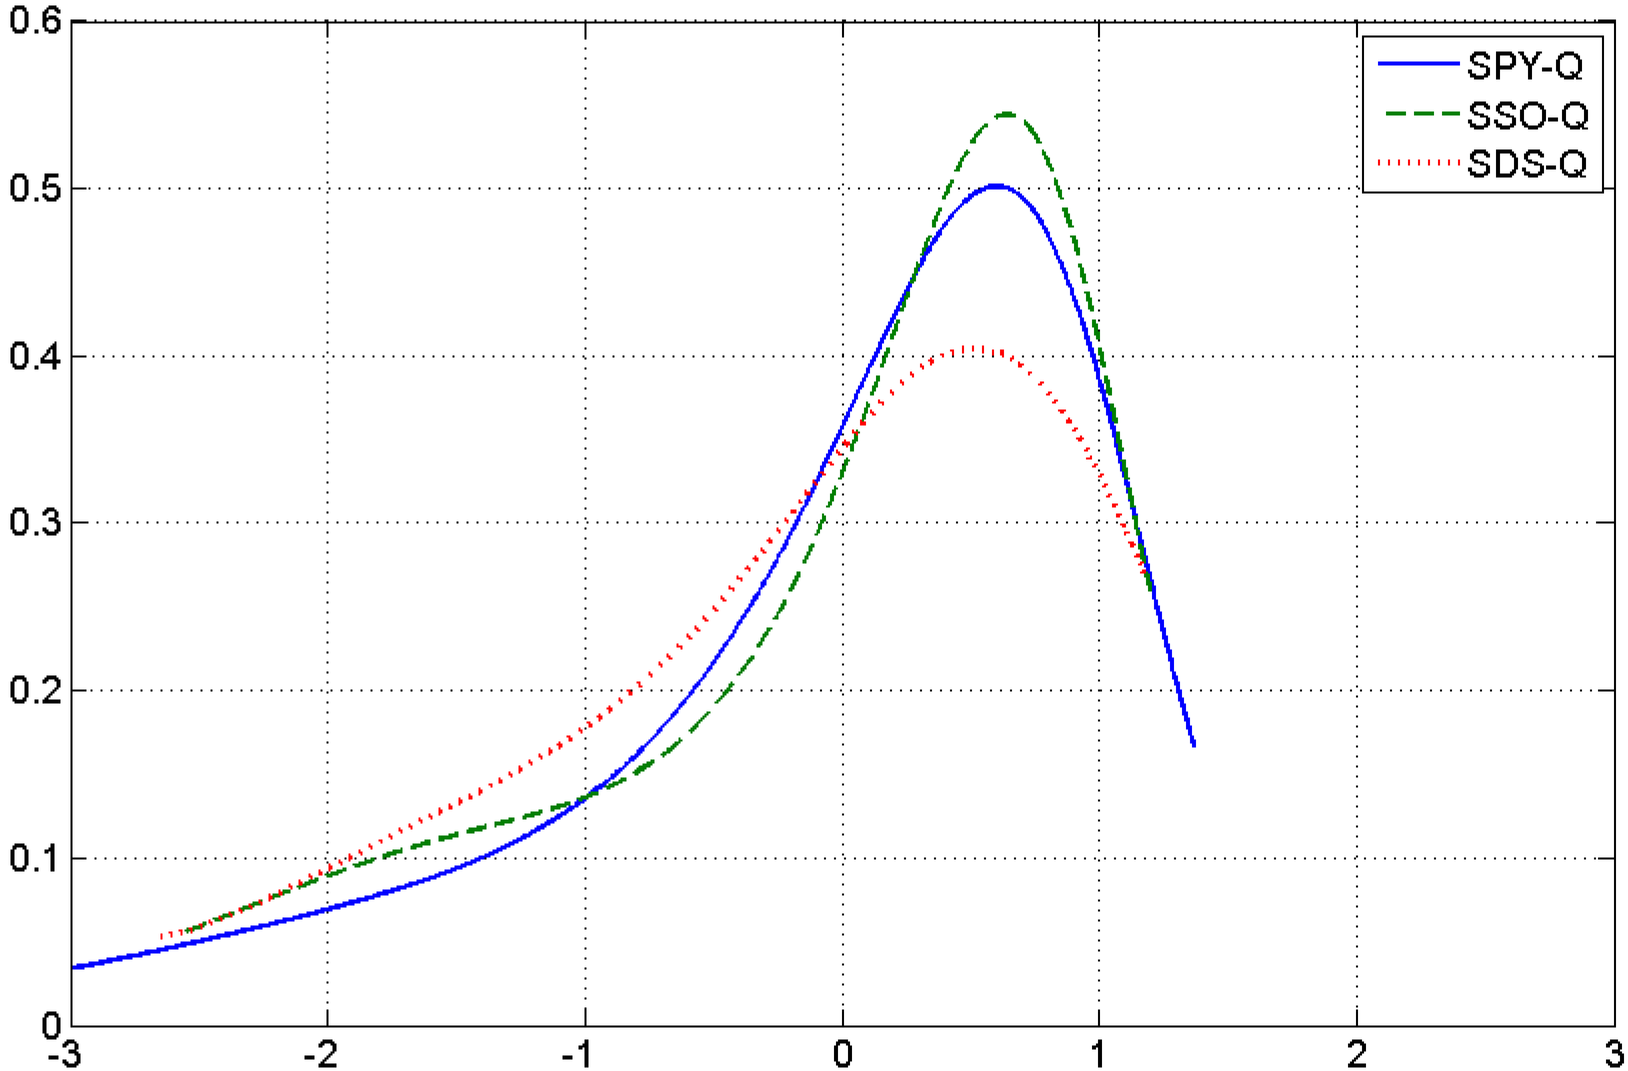

02-May-2008; T = 50; SPX = 1413.90; dSPX = 0.32%; VIX = 18.18; dVIX = -0.70; p-dist'n vol'y = 15.75

Figure 3: Average Relative Demand Intensities in the Full Sample

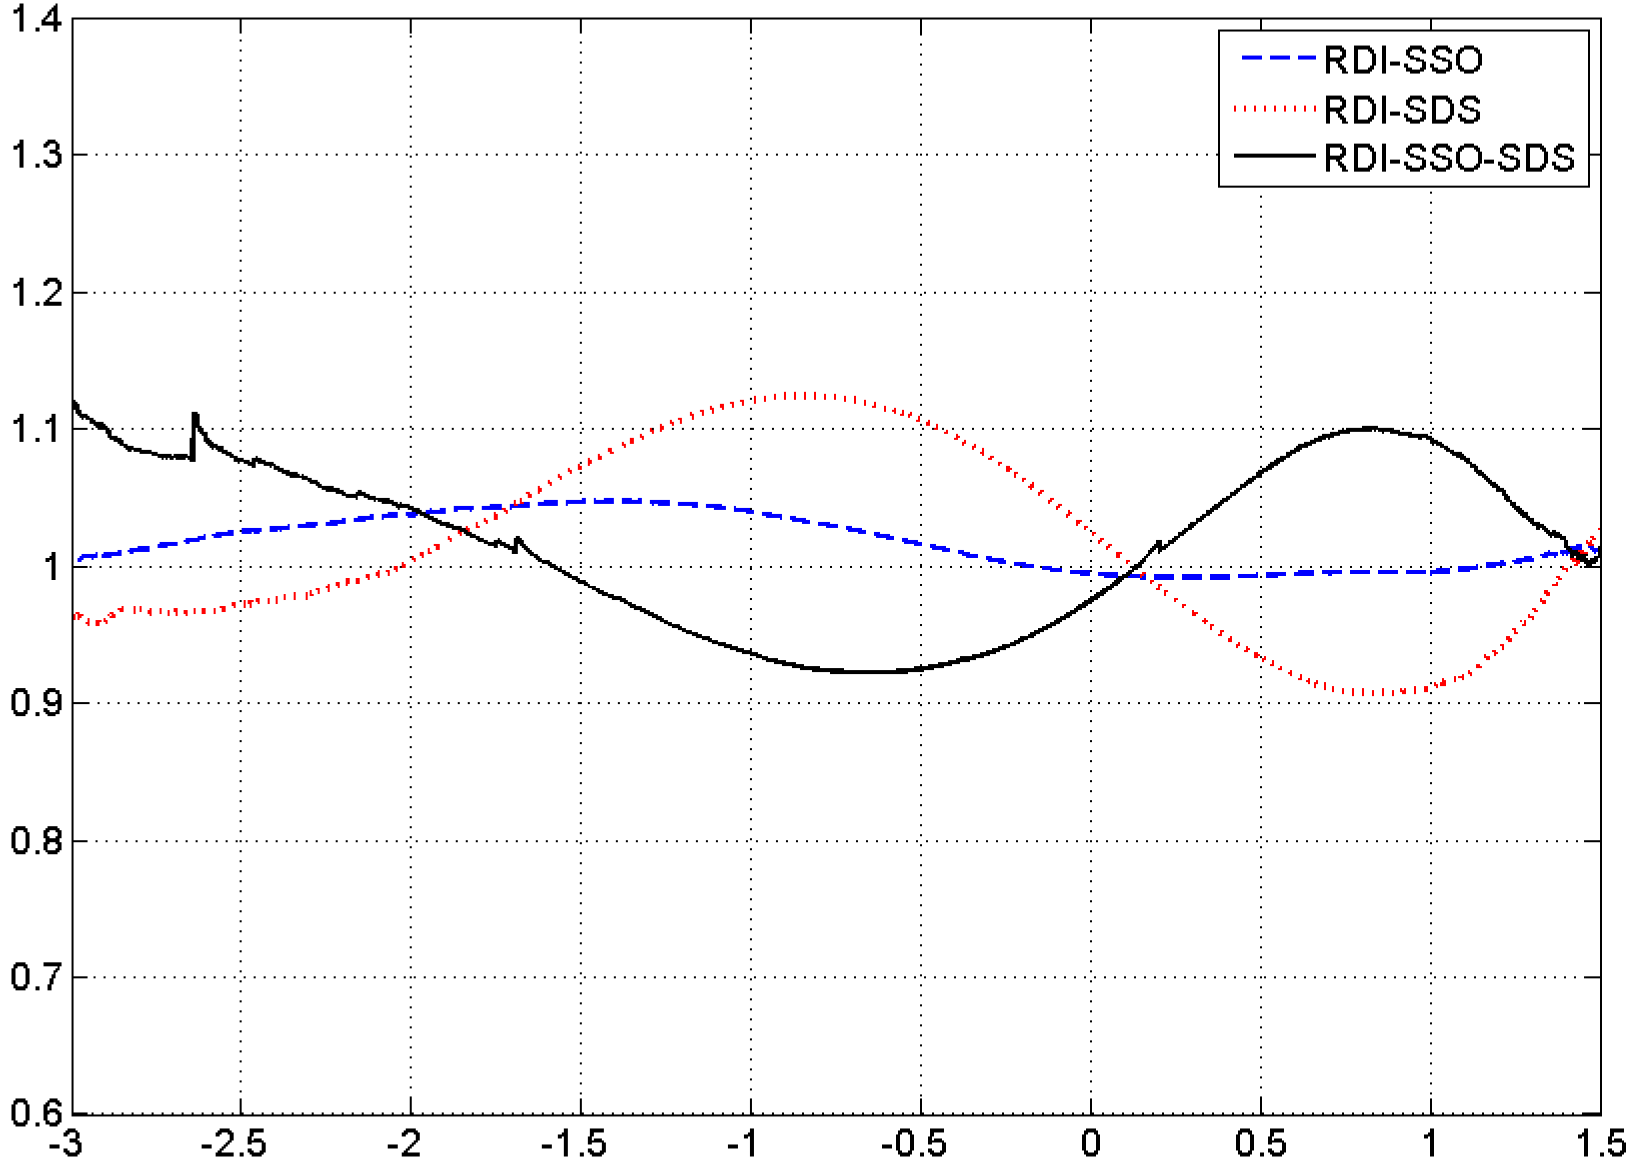

Figure 4: Pricing Kernels on May 2, 2008 Assuming Homogeneous Beliefs on p-density

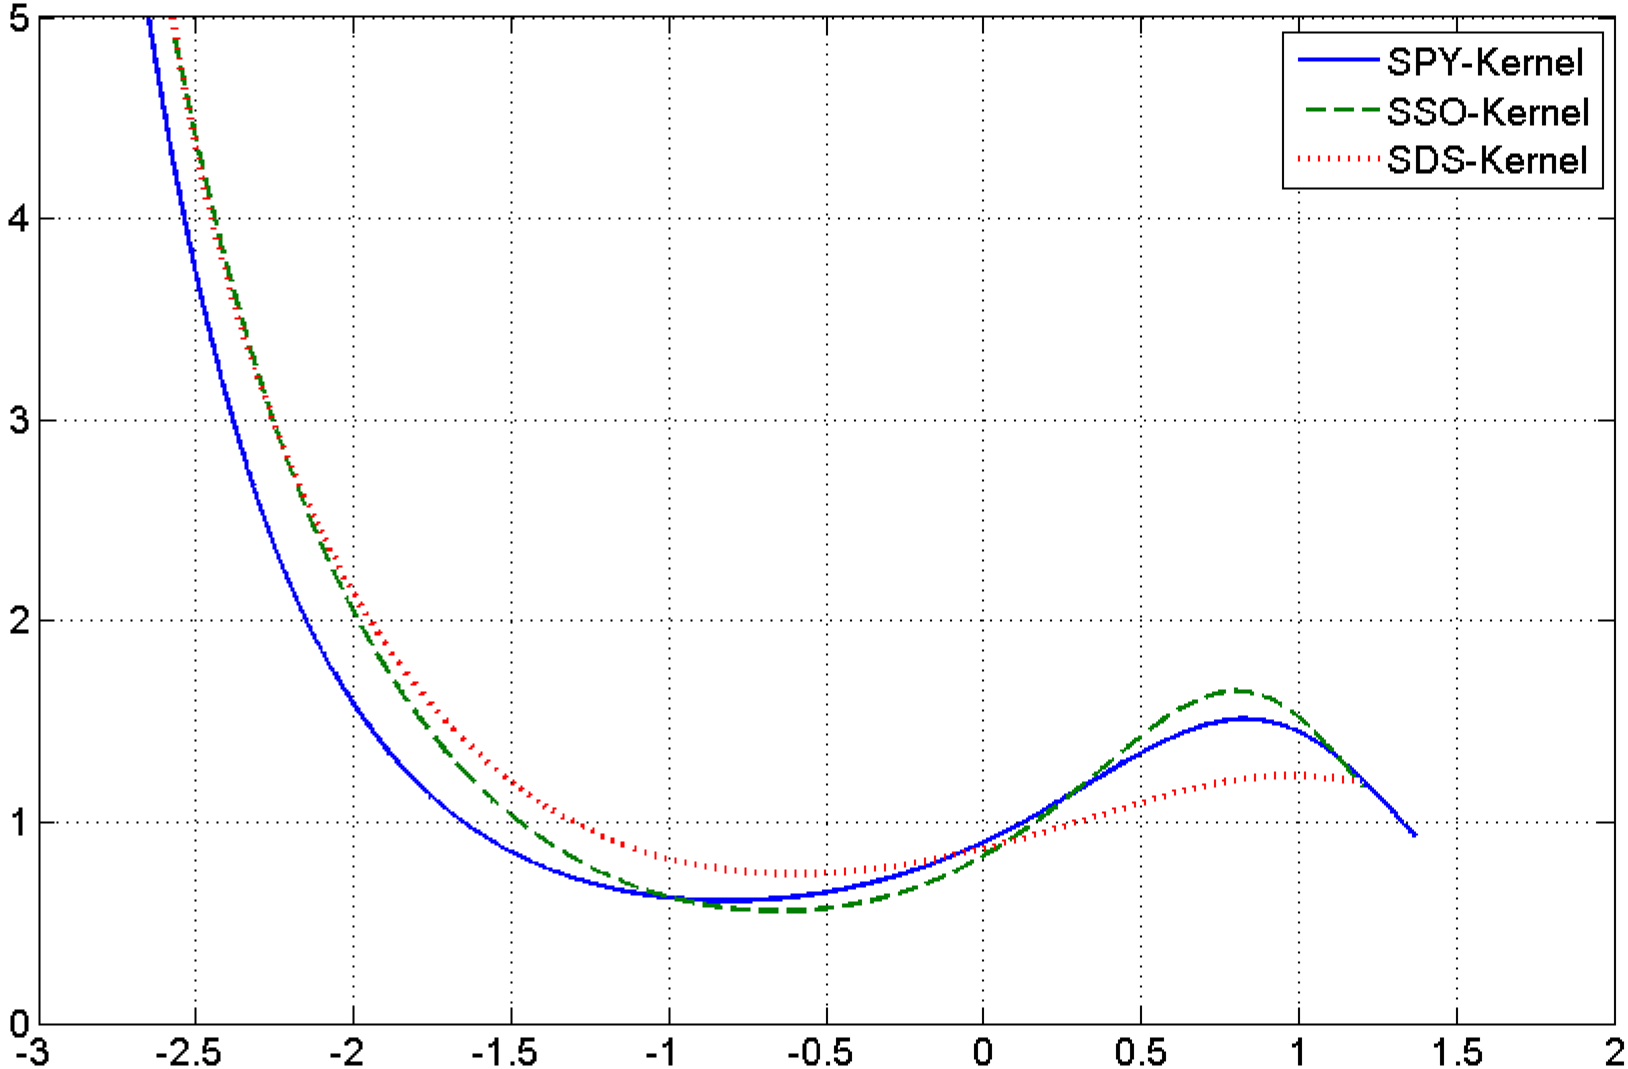

Figure 5: Implied p-densities for May 2, 2008 Assuming Homogeneous Risk Preferences