NYU Stern – Federal Reserve Conference on Risk Neutral Densities

Some Results on Extracting and Understanding The Risk Neutral Returns Distribution for the U.S. Stock Market

Stephen Figlewski**

This presentation includes early results from a joint research project with Muhammad Fahd Malik, a student in Financial Mathematics at NYU's Courant Institute.

** Professor of Finance
New York University Stern School of Business
44 West 4th Street
New York, NY 10012

sfiglews@stern.nyu.edu

Thanks to NASDAQ OMX for their generous financial support of the
NYU Stern School Derivatives Research Project

Research Overview

Figlewski. "Estimating the Implied Risk Neutral Density for the U.S. Market Portfolio,"
in Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle, Oxford, U.K.: Oxford University Press, 2010.

- RND is always left skewed
- The left tail responds more than the right to the market return

"The Impact of the Federal Reserve's Interest Rate Target Announcement on Stock Prices: A Closer Look at How the Market Impounds New Information."
with Justin Birru. Working paper 2010, in revision.

- Uses the RND in real time to explore the impact and reverberations of the Fed funds target announcement in the stock market
- The reduction in risk neutral VAR(S_T) can be used as a measure of information content

"What is Risk Neutral Volatility?" Working paper 2013.

- explores importance of different volatility-related measures (e.g. GARCH vs historical volatility; trading range)
- also explores "sentiment" measures (e.g., Michigan Survey of Consumer Sentiment)

Extracting the Risk Neutral Density from Options Prices

Breeden and Litzenberger (Journal of Business, 1978) showed how the risk neutral probability distribution for S_T, the value of the underlying asset on option expiration day can be extracted from a set of market option prices without a pricing model.

Two major problems in constructing a complete risk neutral density from a set of market option prices are:

1. The procedure uses market option prices, but exercise prices for traded option contracts may be far apart, so to produce a smooth density one must smooth and interpolate option prices to limit pricing noise.

2. Only a limited range of exercise prices is traded, so some way to extend the distribution to the tails is needed.

But one terrific advantage is that, unlike implied volatility, the risk neutral density is model-free.

Extracting the Risk Neutral Density from Options Prices

The value of a call option is the expected value under the risk neutral distribution of its payoff on the expiration date T, discounted back to the present.

$$C = e^{-rT} \int_X^\infty (S_T - X) f(S_T) dS_T$$

Taking the partial derivative with respect to X,

$$\frac{\partial C}{\partial X} = -e^{-rT} \int_X^\infty f(S_T) dS_T = -e^{-rT} \left[1 - F(X)\right]$$

$$\boxed{F(X) = e^{rT} \frac{\partial C}{\partial X} + 1}$$

F(X) is the cumulative probability distribution. Taking the derivative again gives f(X), the risk neutral probability density, RND.

$$\boxed{f(X) = e^{rT} \frac{\partial^2 C}{\partial X^2}}$$

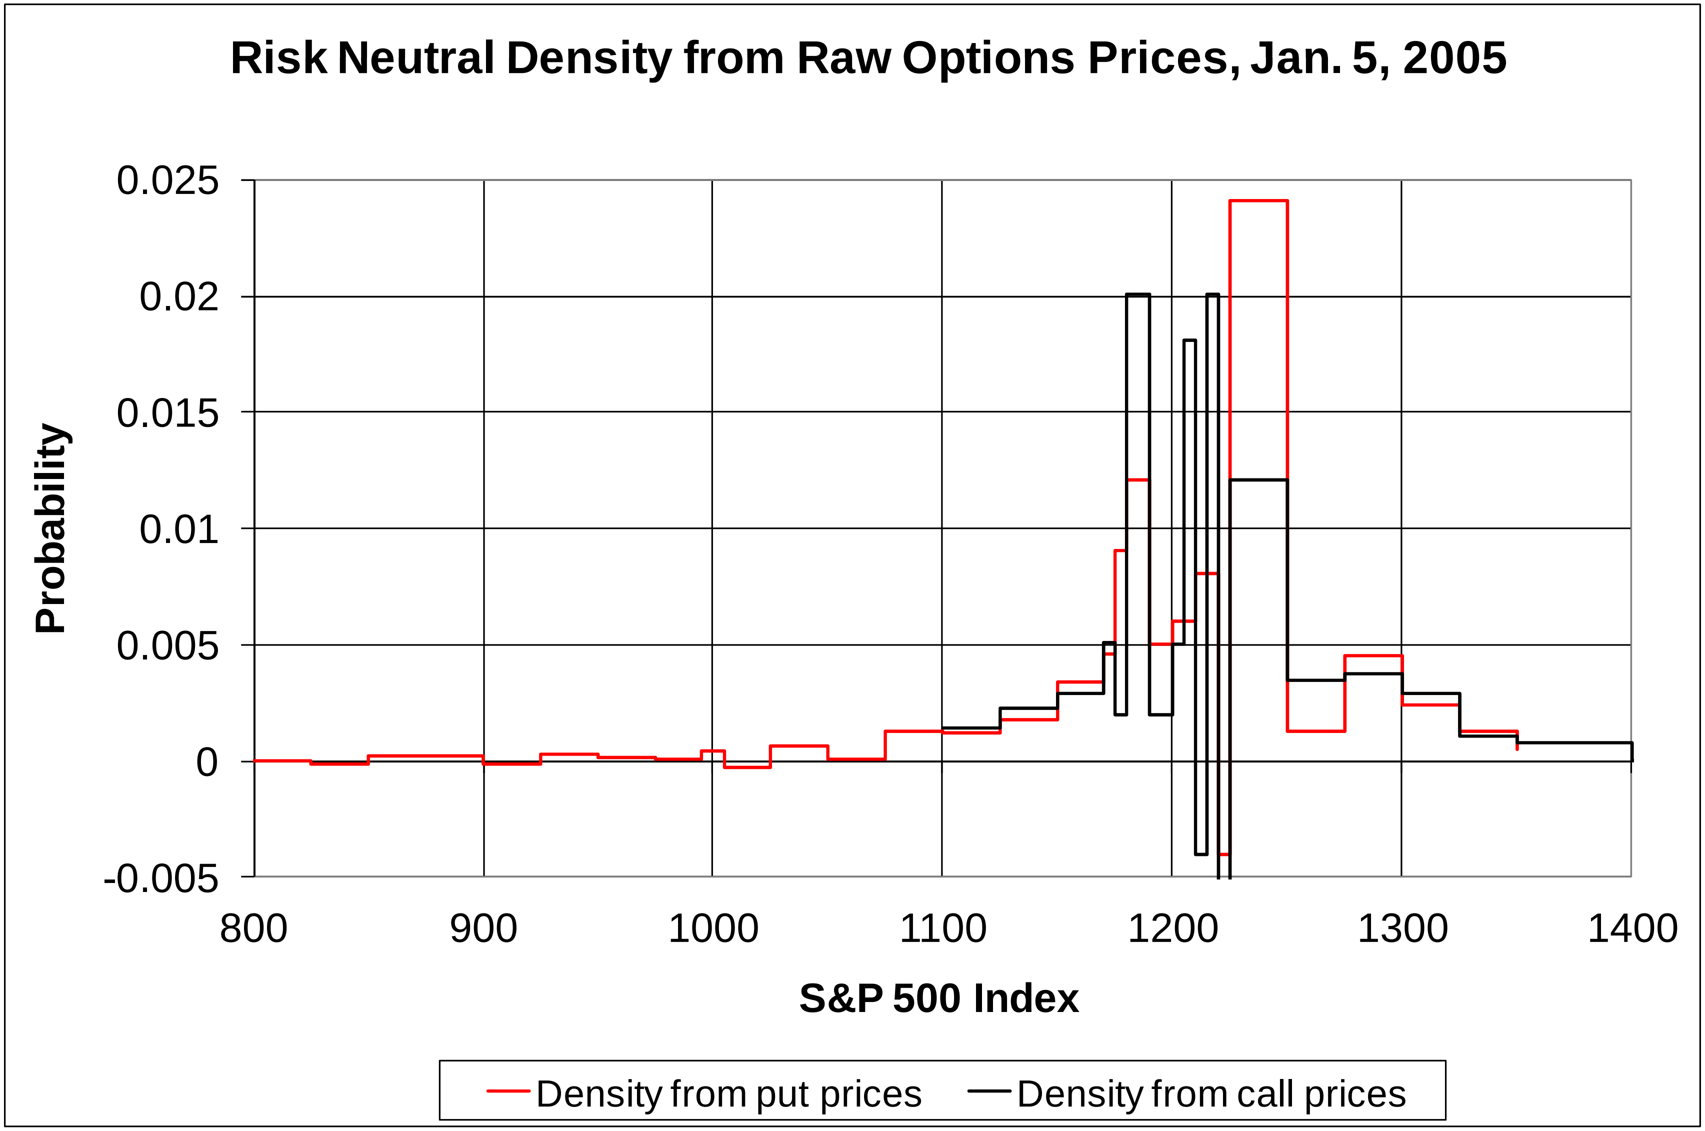
Risk Neutral Density from Raw Options Prices, Jan. 5, 2005
0.025
0.02
0.015
0.01
0.005
0
-0.005
Probability
800
900
1000
1100
1200
1300
1400
S&P 500 Index
Density from put prices
Density from call prices

Extracting the Risk Neutral Density from Options Prices in Practice

Obtaining a well-behaved risk neutral density from market option prices is a nontrivial exercise. Here are the main steps we follow.

1. Use bid and ask quotes, rather than transactions prices. Eliminate options too far in or out of the money.

2. Use out of the money calls, out of the money puts, and a blend of the two at the money

3. Convert prices to Black-Scholes implied volatilities

4. Interpolate the IVs using a 4th degree smoothing spline fitted to the IV bid-ask spread

5. Convert the interpolated IV curve back to option prices and extract the middle portion of the risk neutral density, using the Breeden-Litzenberger strategy

6. Append tails to the Risk Neutral Density from a Generalized Pareto Distribution (GPD)

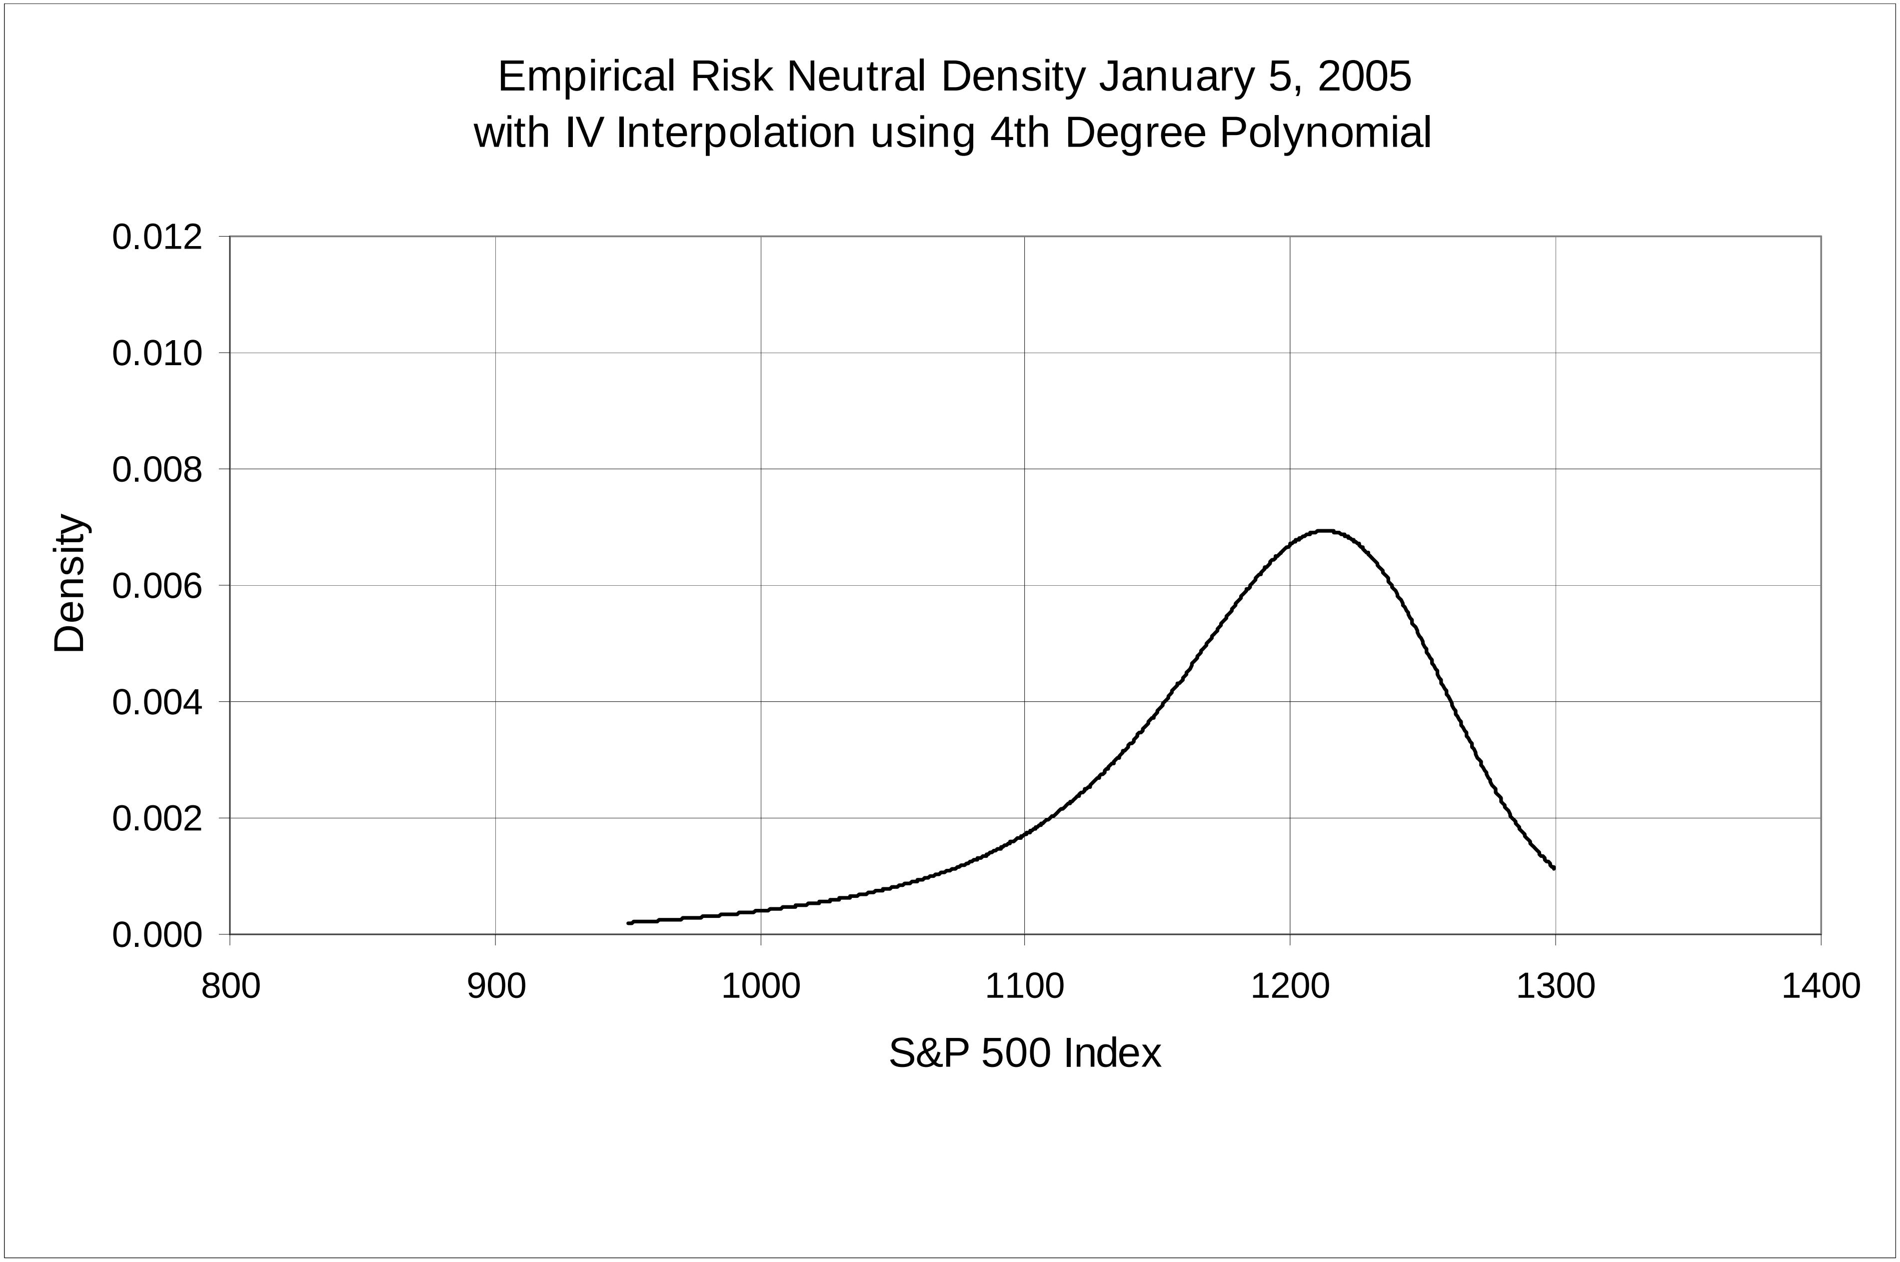
Empirical Risk Neutral Density January 5, 2005
with IV Interpolation using 4th Degree Polynomial
Density
0.012
0.010
0.008
0.006
0.004
0.002
0.000
800
900
1000
1100
1200
1300
1400
S&P 500 Index

Extending the Risk Neutral Density into the Tails

Only the middle portion of the empirical RND spanned by the set of option strikes can be estimated this way.

Some estimation techniques generate tails automatically, like fitting one or a mixture of parametric distributions (e.g., lognormal, Student-t) to option prices.

We take the empirical RND from the data and append tails from a Generalized Pareto (GPD) distribution. The fitted tails contain the correct total probability and are fitted to match the shape of the extreme portions of the empirical RND.

By the Fisher-Tippett Theorem, the tails of any member of a very broad class of distributions in the limit resemble the tail of a Generalized Extreme Value distribution.

The GPD is the distribution of random draws from a density function, that exceed some high threshold value. It is closely related to the Generalized Extreme Value (GEV) distribution. In particular, it has the same tail shape parameter. Either density works well for our purpose, but the GPD has the advantage of only requiring estimation of two parameters, rather than 3 for the GEV.

Birru and Figlewski (JFM 2012) show that the GEV and GPD produce much more accurate tails than the lognormal.

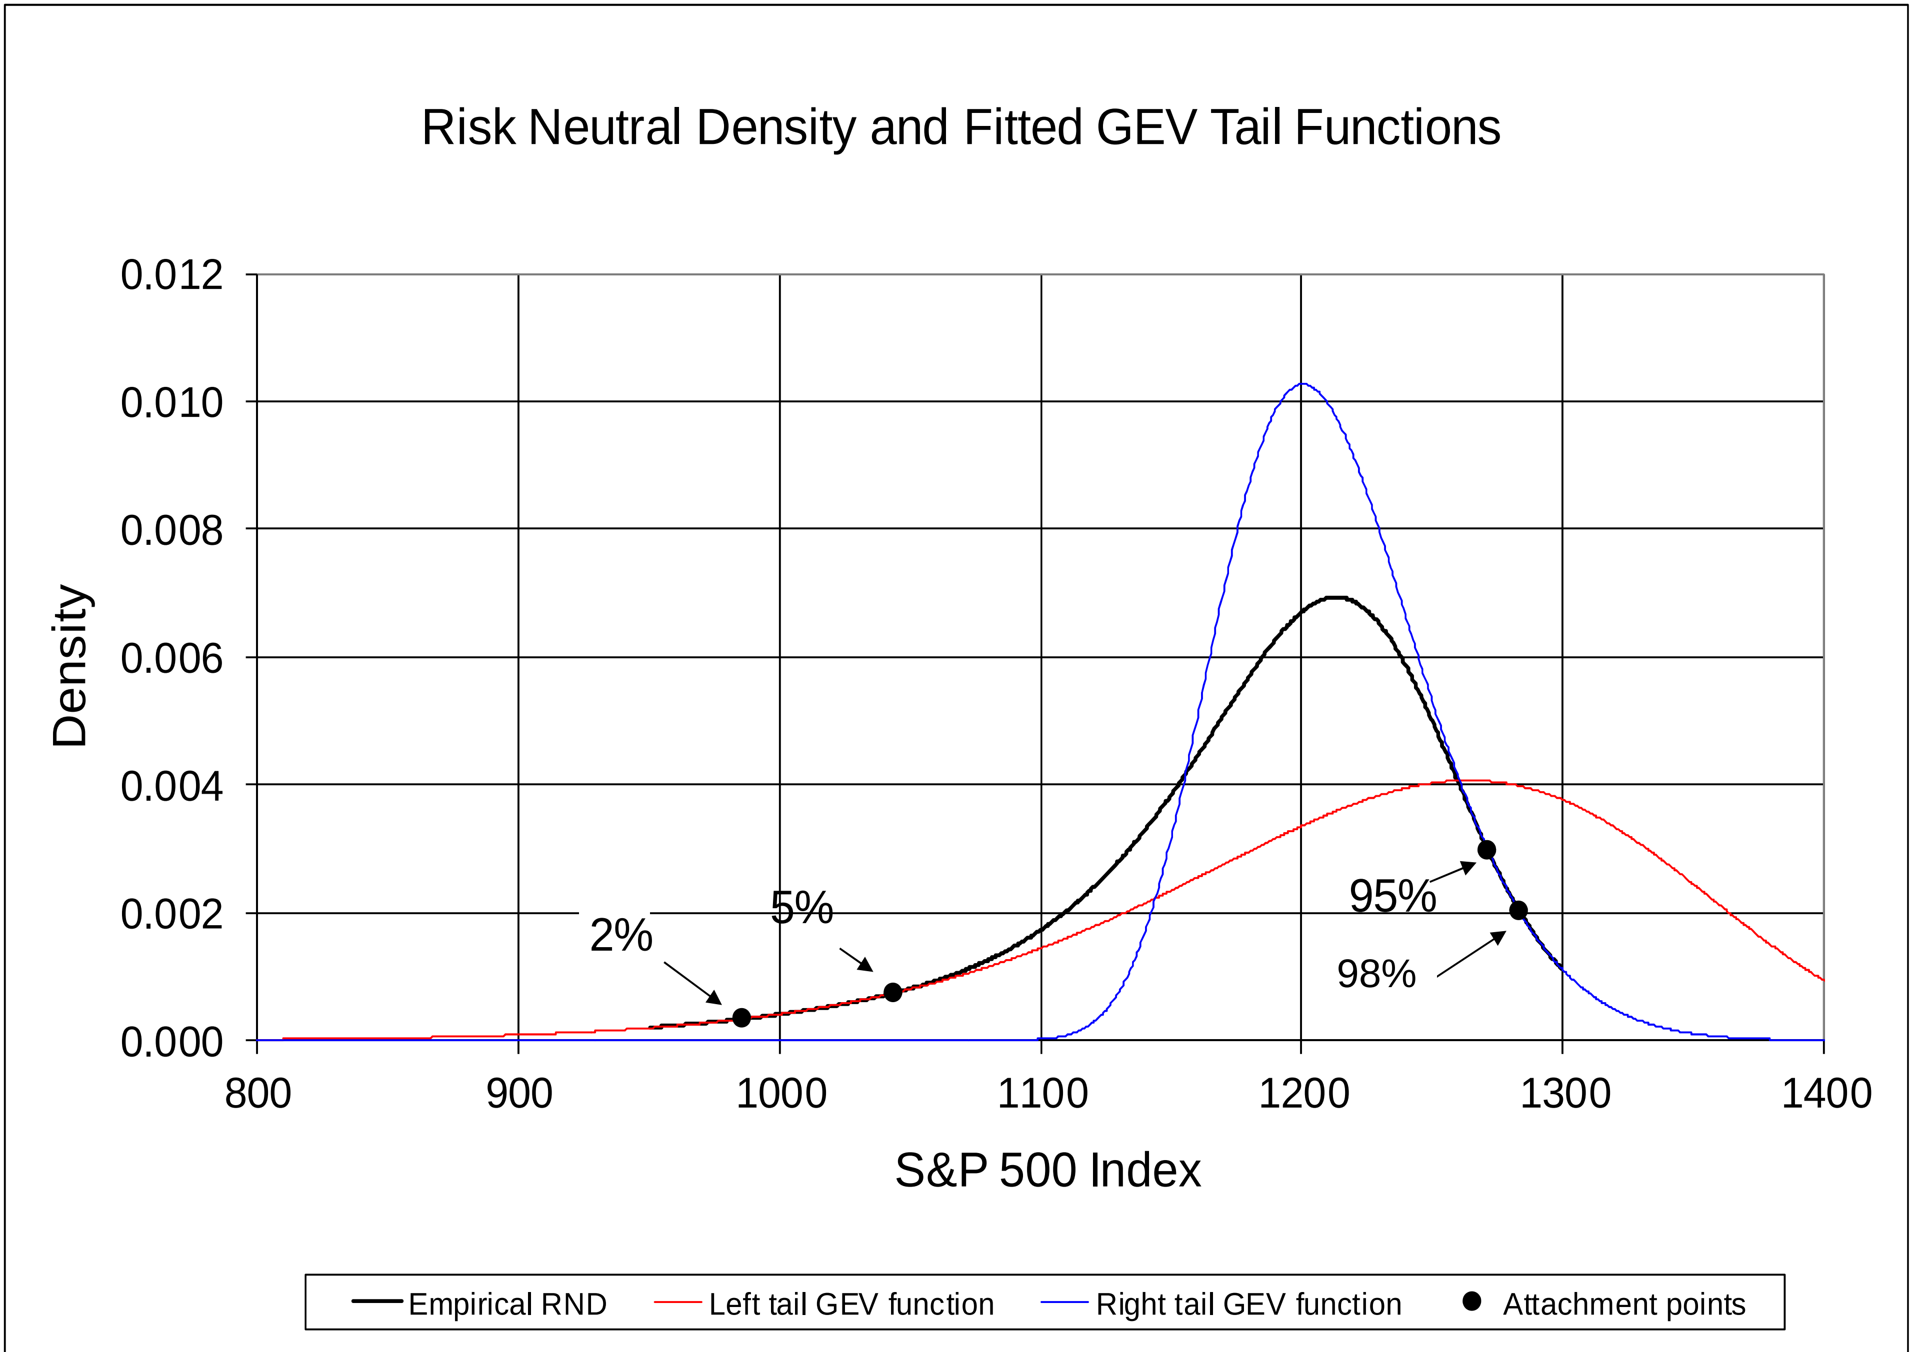
Risk Neutral Density and Fitted GEV Tail Functions
Density
0.012
0.010
0.008
0.006
0.004
0.002
0.000
800
900
1000
1100
1200
1300
1400
S&P 500 Index
2%
5%
95%
98%
Empirical RND
Left tail GEV function
Right tail GEV function
Attachment points

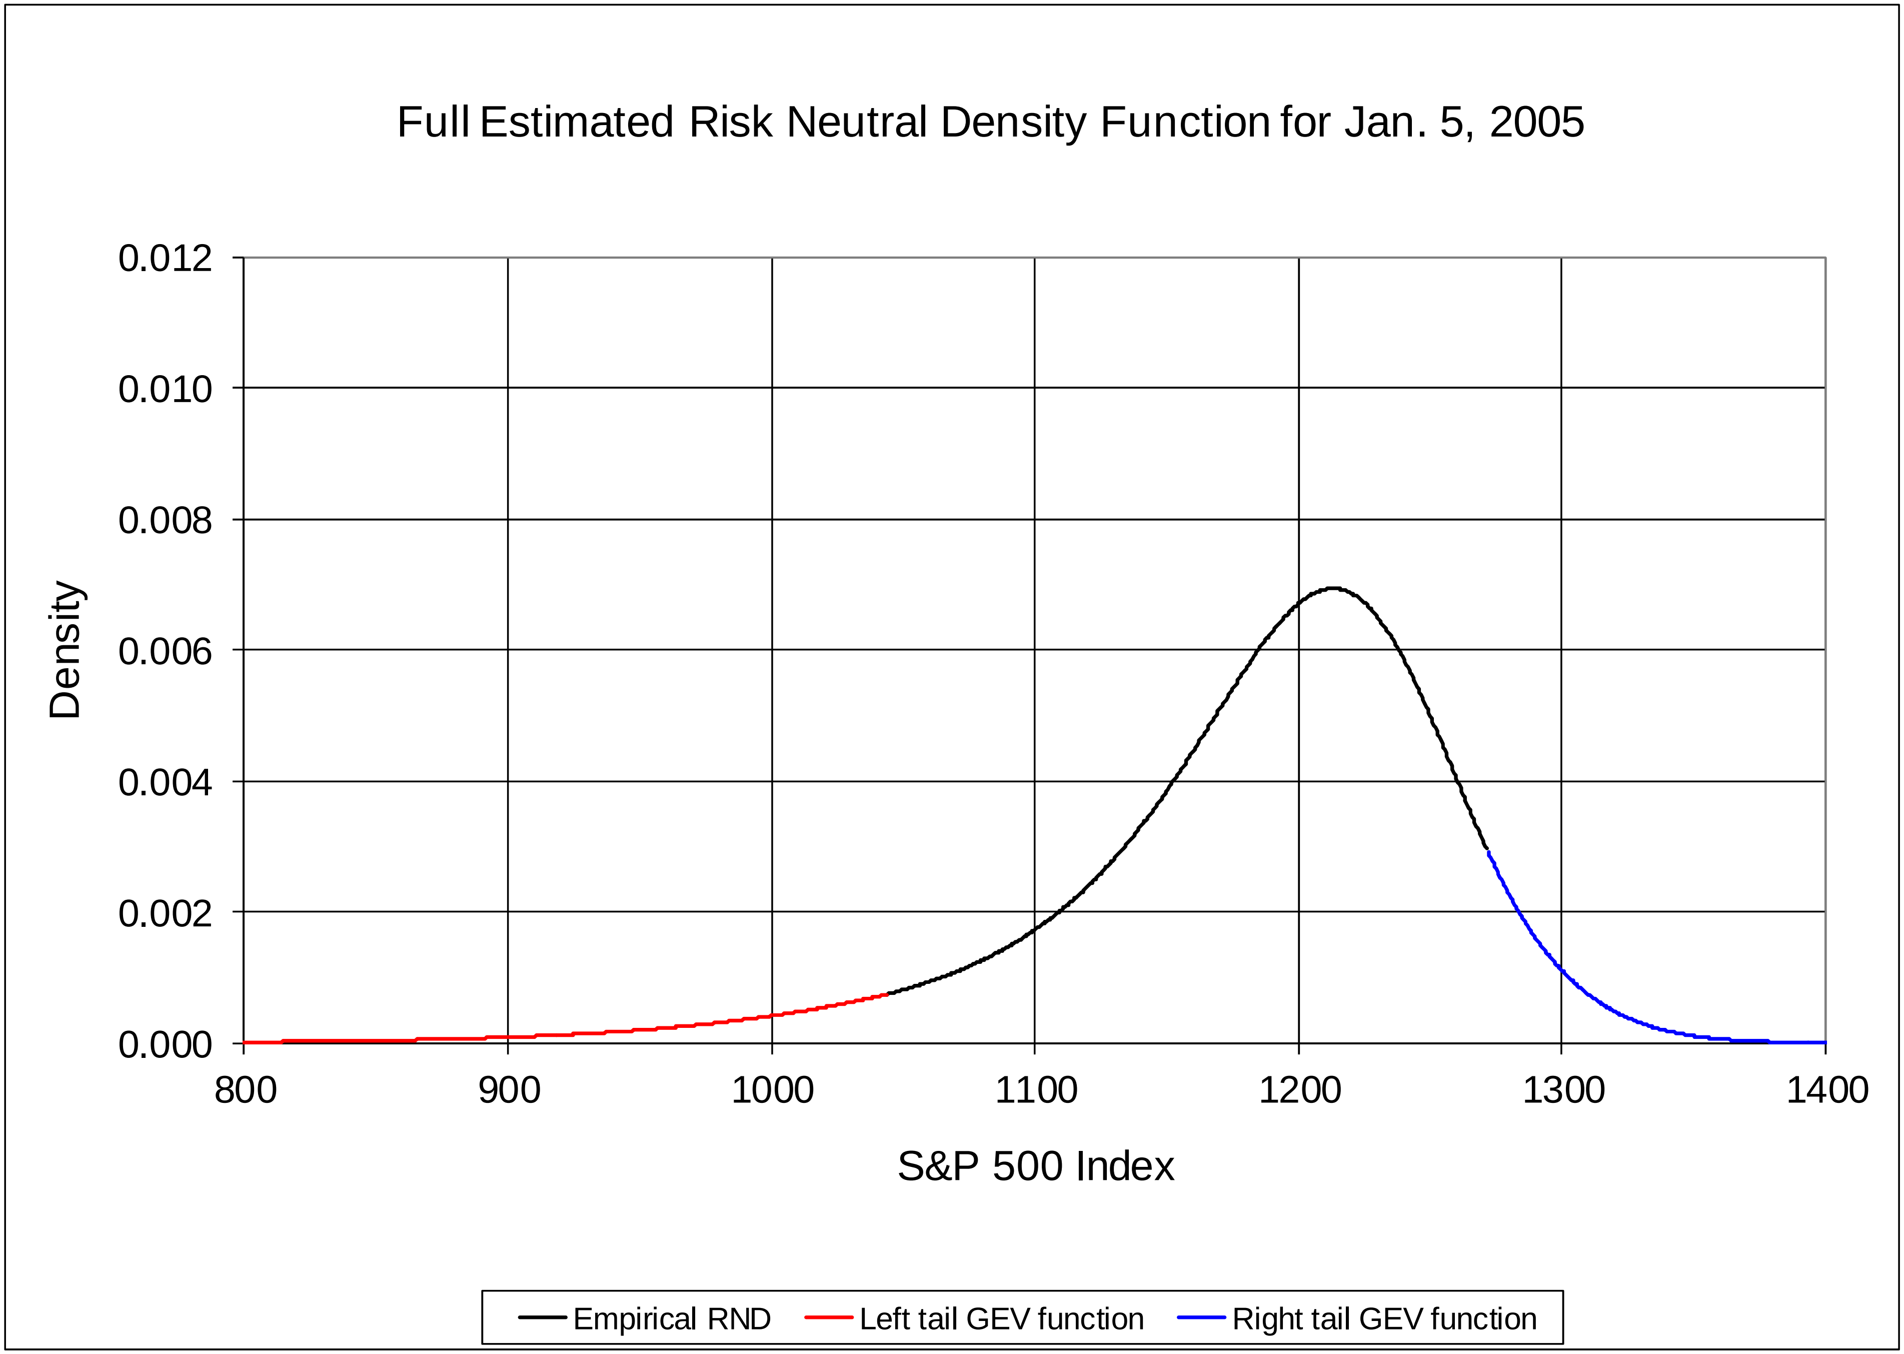

Full Estimated Risk Neutral Density Function for Jan. 5, 2005
Density
0.012
0.010
0.008
0.006
0.004
0.002
0.000
800
900
1000
1100
1200
1300
1400
S&P 500 Index
Empirical RND
Left tail GEV function
Right tail GEV function

Some Risk Neutral Densities

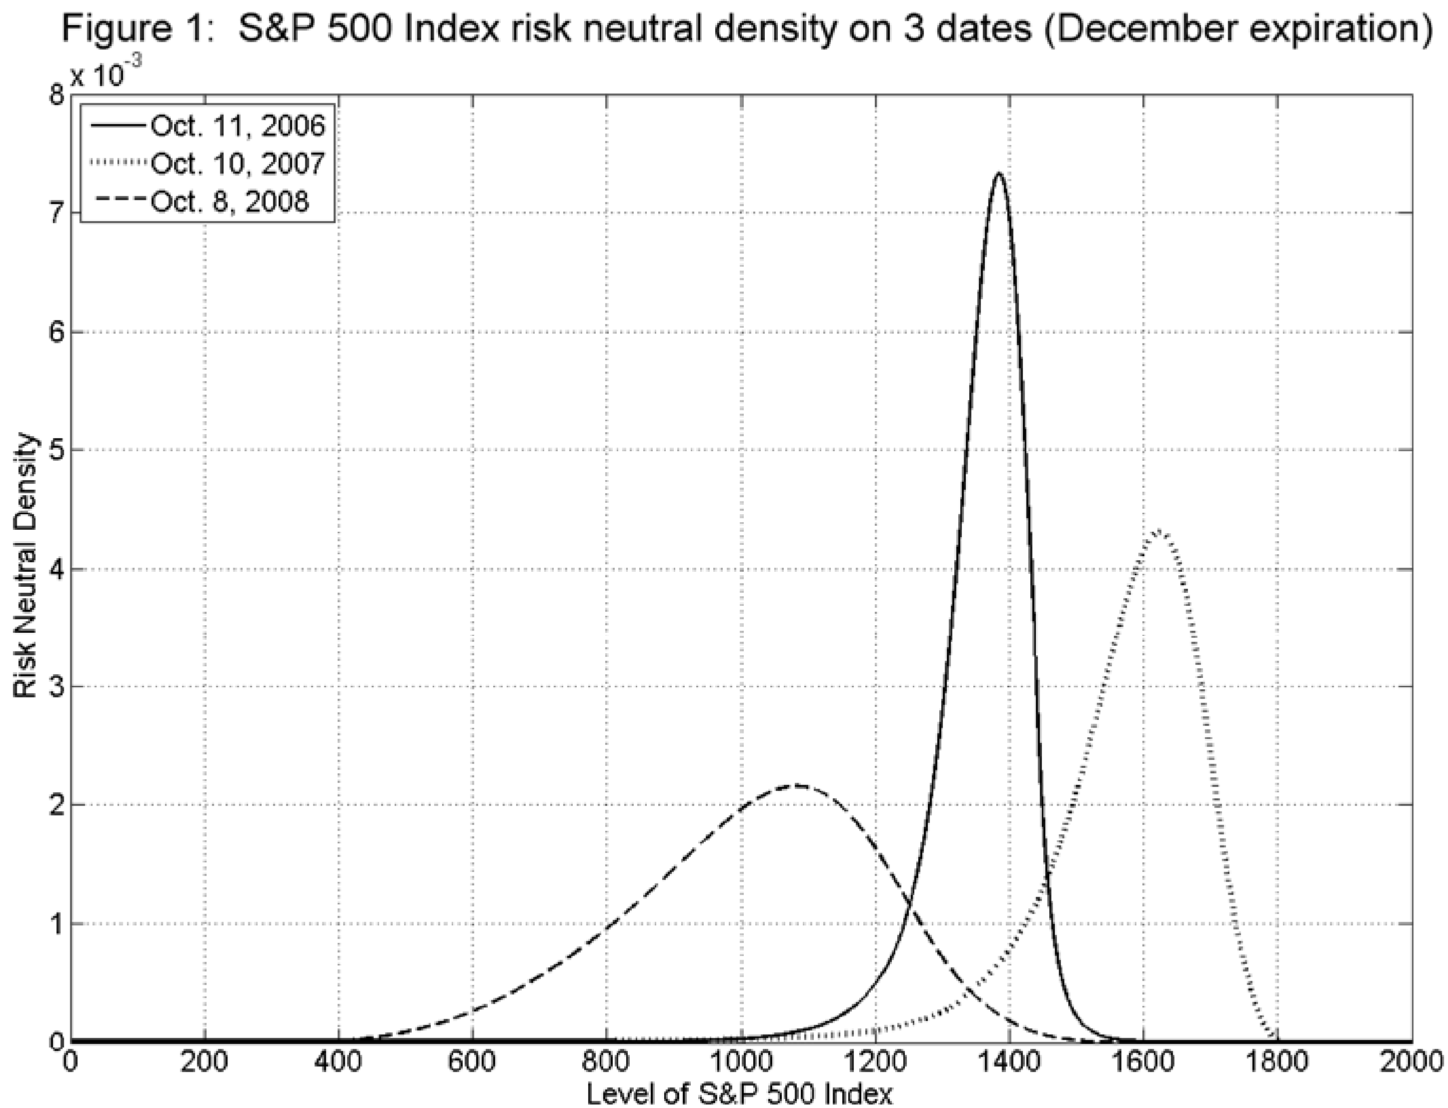

Exchange-Traded Funds Based on the S&P 500 Index

An Exchange Traded Fund (ETF) is a stock portfolio, similar to a closed-end mutual fund, typically designed to provide focused exposure to a specific market factor.

The (original) SPDR contract (ticker symbol SPY) attempts to earn the same return each day as on the "market portfolio" represented by the S&P 500 index.

Other ETFs tied to the same S&P 500 index try to duplicate the return on a leveraged investment in the index.

- The SSO "double long" ETF attempts to earn twice the daily return on the S&P 500 index portfolio.
- The SDS "double short" ETF tries to match twice the return on a short position in the index.

Example: If the S&P 500 index portfolio return is 1.0% on date t,

- the SPY should earn 1.0%,
- the SSO return should be 2.0%
- the SDS return should be -2.0%.

Heterogeneity and the Representative Agent

A risk neutral density extracted from SPY options reflects the probability distribution investors currently expect for the S&P 500 index on option expiration day, modified by their required risk premia.

Individual investors have heterogeneous risk preferences and probability estimates, which the market must aggregate in forming a single market price for each security.

Financial models often avoid the aggregation problem by treating investors as being identical, equal to some "representative agent."

In a market that can be modeled as having a single representative agent, security prices can be directly connected to that agent's utility function.

The Stochastic Discount Factor (or Pricing Kernel)

In pricing a security, every possible payoff it might have in the future is weighted by its probability, modified to incorporate risk premia.

An option on the SPY ETF will reflect the representative agent's estimate of the true probability distribution for the future level of the whole stock market, "risk-neutralized" based on how desirable a dollar of payoff is in each case.

Let $p(S_T)$ be what investors believe to be the true probability density for the level of the S&P index on option expiration date T, and $q(S_T)$ be the risk neutral density (RND). Risk-neutralization is often expressed in terms of the "stochastic discount factor" or "pricing kernel" that connects the "p" and "q" distributions. The pricing kernel, call it $m(S_T)$, gives the value today of a dollar payoff in each possible future "state of the world" S_T.

The market price for a security today is the expected value of the future payoff in dollars multiplied by the pricing kernel $m(S_T)$

$$C_0 \;=\; E\,[payoff(S_T) \times m(S_T)]$$

The Stochastic Discount Factor (or Pricing Kernel)

For an SPY call option,

$$C_0 = E[payoff(S_T) \times m(S_T)]$$

In terms of the RND, $q(S_T)$:

$$C_0 = e^{-rT} \int_X^{\infty} (S_T - X) q(S_T) dS_T$$

$$= e^{-rT} \int_X^{\infty} (S_T - X) p(S_T) \left({q(S_T)}/{p(S_T)} \right) dS_T$$

This expression shows the future payoffs weighted by the true probabilities $p(S_T)$, and then multiplied by the risk neutralization factor $q(S_T)/p(S_T)$.

If investors really were risk neutral, q(.) would be the same as p(.), the ratio would be 1.0 for all values of S_T, and the option price would be its true expected payoff discounted back to the present at the riskless interest rate r.

The Stochastic Discount Factor (or Pricing Kernel)

The stock market portfolio is often taken to be a proxy for total wealth, so a dollar should have greater utility when the stock market has fallen (overall wealth is low) than when it has gone up (higher total wealth).

In that case, we should have $q(S_T) > p(S_T)$ for low S_T and $q(S_T) < p(S_T)$ for high S_T .

In general, risk neutralization has the effect of increasing the effective probability for states of the world that investors don't like (low stock prices, in this case) and lowering the effective probability for states that investors do like (high stock prices).

Is this what happens in the real world? Let's see.

SPY Option Prices and the Pricing Kernel

The risk neutral q(.) is extracted from SPY options. To compute the stochastic discount factor, we also need an estimated "true" density p(.).

Assume that on date t the market believes the probability density for the level of the S&P index on date T is lognormal.

- The expected mean return is the current riskless interest rate plus an annualized risk premium of 5.0% (5% is the about average risk premium used by finance professors, according to a recent survey).
- The expected volatility is a combination of historical volatility over the previous 63 trading days up to the present date t and the date t level of the VIX implied volatility index.

The weights for each date t come from running this regression on all daily data from Jan. 2, 1990 up to date t, where Realized volatility is computed over a horizon of length (T-t):

$$\text{Realized volatility} = a + b \text{ Historical volatility} + c \text{ VIX} + \varepsilon$$

SPY Option Prices and the Pricing Kernel

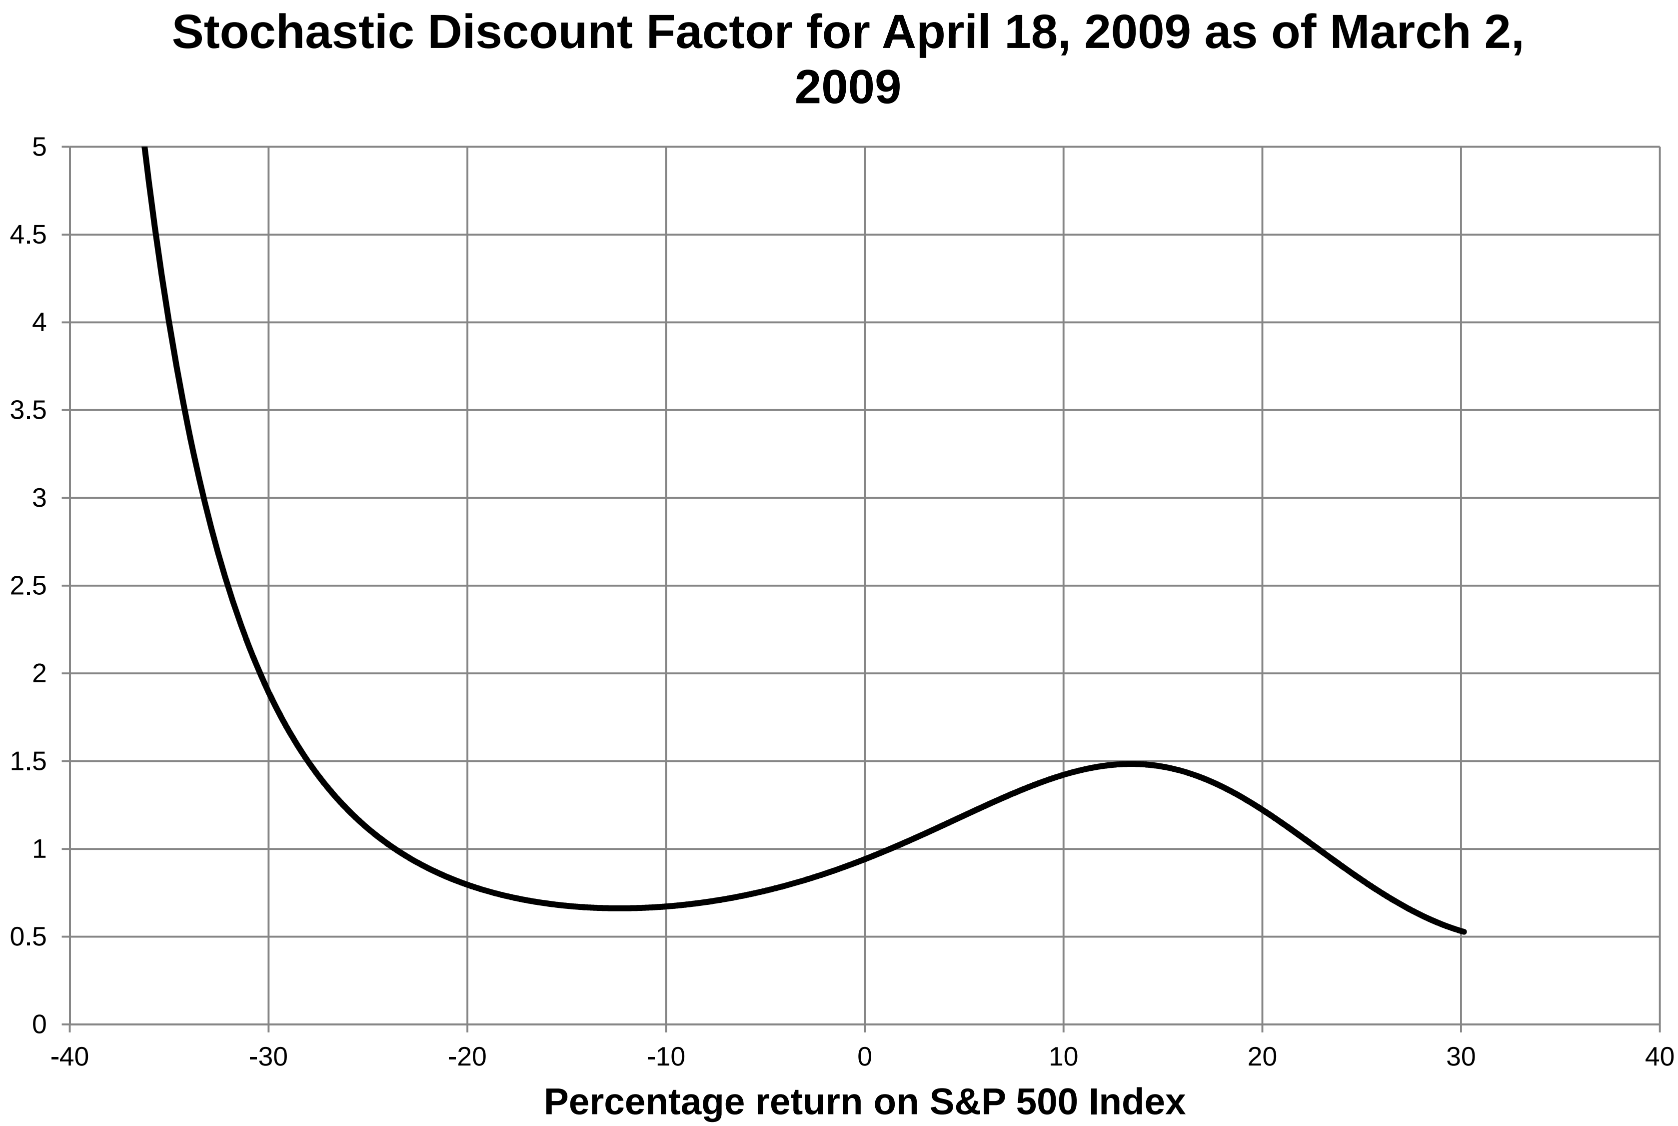

Oops!

The left and right ends of the curve are downward sloping, as expected, but the upward sloping portion in the middle is strange. It says that investors value a dollar of payoff less if the stock market goes down 10% in the next 8 weeks than if stocks go up 10%.

This anomalous pricing behavior is present throughout our data and other researchers have consistently found it for the US stock market. It has come to be known as the "stochastic discount factor puzzle" or the "pricing kernel puzzle."

A Possible Solution to the Pricing Kernel Puzzle

Hersh Shefrin (A Behavioral Approach to Asset Pricing, 2008) offers one of the most appealing explanations for the pricing kernel puzzle. He shows that, rather than assuming a representative agent, simply allowing investors to have different expectations about the mean and the volatility of future returns is enough to produce shapes like this one.

The ratio $q(S_T)/p(S_T)$ gives the intensity of market demand for a dollar of payoff when the S&P index on date T is S_T, relative to how a risk neutral investor would value that dollar.

The ratio for a given S_T is greater than 1.0 if investors' become more averse to owning a risky security when they have less wealth, and investors have less wealth at this S_T. Risk aversion is the standard explanation for the downward sloping portions of the pricing kernel.

But the RND also reflects investors' expectations. An optimistic investor expects the market to rise more rapidly than the average investor does, so he will pay extra for call options with exercise prices above the current market price that are currently out of the money. For him, the lower value of a dollar when S_T is high can be offset by his prediction that S_T is more likely than the rest of the market thinks it is.

Heterogeneous beliefs can produce bumps in the pricing kernel curve.

A Possible Solution to the Pricing Kernel Puzzle

We believe real world investors are heterogeneous in their risk attitudes and especially in their expectations about security returns, and limits to arbitrage allow some pricing discrepancies to persist even among closely related markets.

In that case, related financial instruments can appeal to different clienteles of investors.

- The q-density may be above the p-density at some S_T because this is an especially unfavorable outcome for the security's clientele and a payoff at that S_T is a hedge.
- The q-density may be above the p-density for any S_T that its clientele believes is more likely than the rest of the market thinks.

Thus, the ratio of the q-density to the p-density measures the relative intensity of demand for a payoff in a state of the world S_T, between investors whose RND is being measured and a risk neutral investor with expectations equal to the market average prediction of the true probabilities.

We will call this kind of ratio a "Relative Demand Intensity" (RDI).

We can also compute the relative demand intensity of one security's clientele relative to another security's clientele by taking the ratio of the two q-densities.

Relative Demand Intensity among Leveraged ETFs

The investors in the three "leveraged" ETF securities, SPY, SSO, and SDS, can be expected to have different expectations about the future returns on the S&P 500 index and possibly different risk preferences also.

We explore those differences using the RNDs extracted from their traded options.

Each ETF's options are written in terms of the price of the underlying ETF. The first step is to put them all on the same basis. We convert the ETF risk neutral densities into implied RNDs for the level of the S&P 500 index on option expiration day.

This is not straightforward, because the multiperiod return from matching a multiple of the S&P return each day is path-dependent. (For example, if S&P return is +1% on date t and -1% on date t+1, the $SSO_{t+2} < SSO_t$.)

Converting from ETF Level to Equivalent S&P 500 Level

Let L be the leverage factor for a given ETF (L_{SSO} = 2; L_{SDS} = -2). S is the S&P500 index and E_t is the price on date t of an ETF tied to S_t with leverage factor L.

For a lognormal diffusion, the expected value of E_T given a realized S_T on a future date T can be computed as:

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

where E_0 and S_0 are the initial prices for the ETF and the S&P index and σ is the volatility of the S&P.

A density $f_E(.)$ defined over values of E_T can be converted to a density $f_S(.)$ over S_T as follows:

$$f_E(E_T) = f_S(S_T) \Bigg/ \left| \frac{S_0}{L \, E_0^{1/L}} \, e^{(L-1)\frac{\sigma^2}{2}T} \, E_T^{L*(\frac{1}{L} - 1)} \right|$$

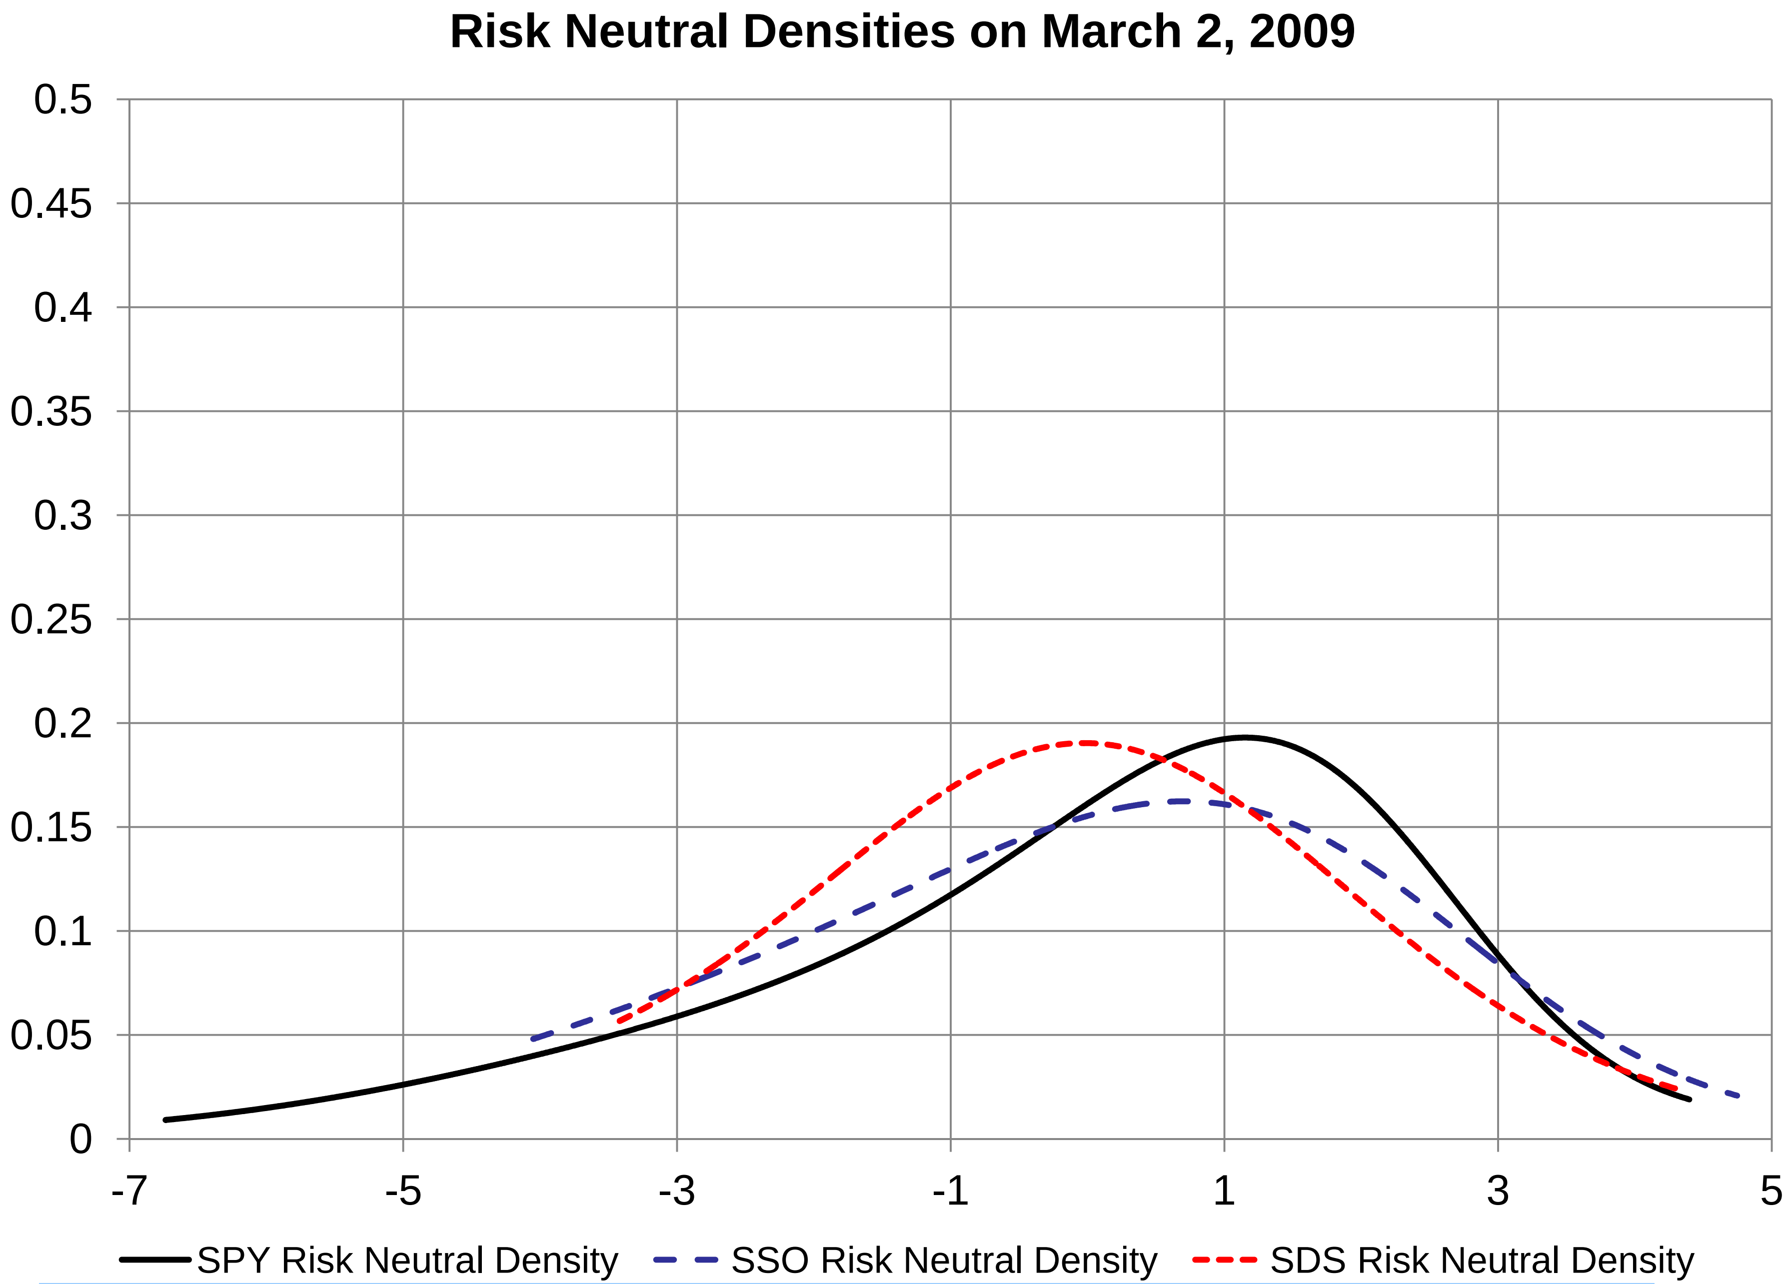
Risk Neutral Densities on March 2, 2009
0.5
0.45
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
-7
-5
-3
-1
1
3
5
SPY Risk Neutral Density
SSO Risk Neutral Density
SDS Risk Neutral Density

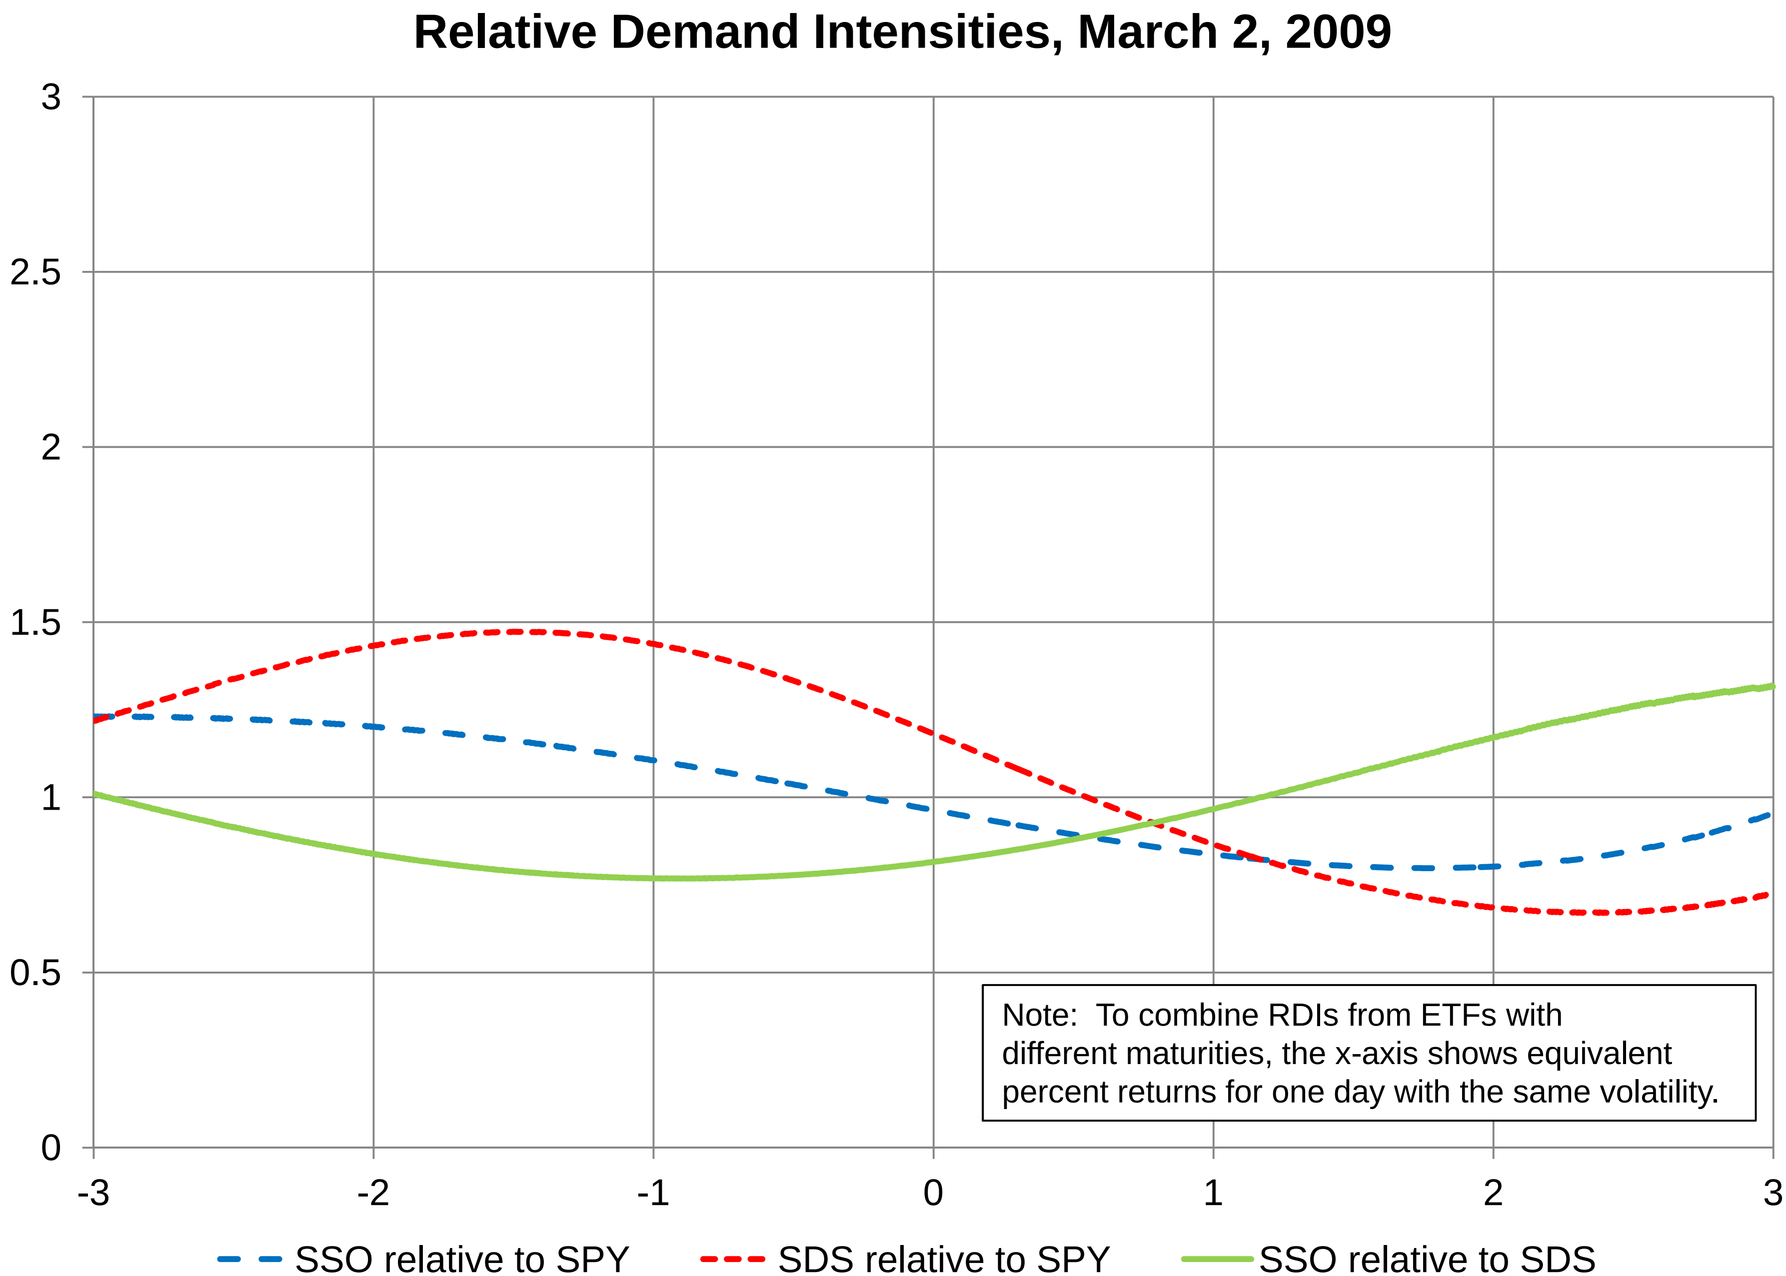
Relative Demand Intensities, March 2, 2009
3
2.5
2
1.5
1
0.5
0
-3
-2
-1
0
1
2
3
Note: To combine RDIs from ETFs with different maturities, the x-axis shows equivalent percent returns for one day with the same volatility.
SSO relative to SPY
SDS relative to SPY
SSO relative to SDS

Average Relative Demand Intensities

	Relative Demand Intensity					
	SSO to SPY		SDS to SPY		SSO to SDS	
1-day Equivalent Return	Mean	Standard deviation	Mean	Standard deviation	Mean	Standard deviation
-3.0%	0.983	0.203	0.943	0.266	1.073	0.504
-2.5%	0.997	0.174	0.978	0.260	1.067	0.502
-2.0%	1.028	0.136	1.014	0.225	1.045	0.219
-1.5%	1.042	0.103	1.058	0.176	0.983	0.186
-1.0%	1.032	0.087	1.111	0.114	0.934	0.084
-0.5%	1.009	0.069	1.102	0.084	0.918	0.069
0.0%	0.989	0.048	1.029	0.061	0.963	0.062
0.5%	0.985	0.063	0.953	0.077	1.036	0.078
1.0%	0.991	0.080	0.937	0.089	1.060	0.086
1.5%	1.017	0.089	0.939	0.101	1.085	0.099
2.0%	1.034	0.105	0.925	0.139	1.123	0.110
2.5%	1.007	0.127	0.822	0.116	1.192	0.094
3.0%	1.070	0.140	0.821	0.084	1.298	0.130

The sample is March 2, 2009 – May 31, 2012; 778 observations.

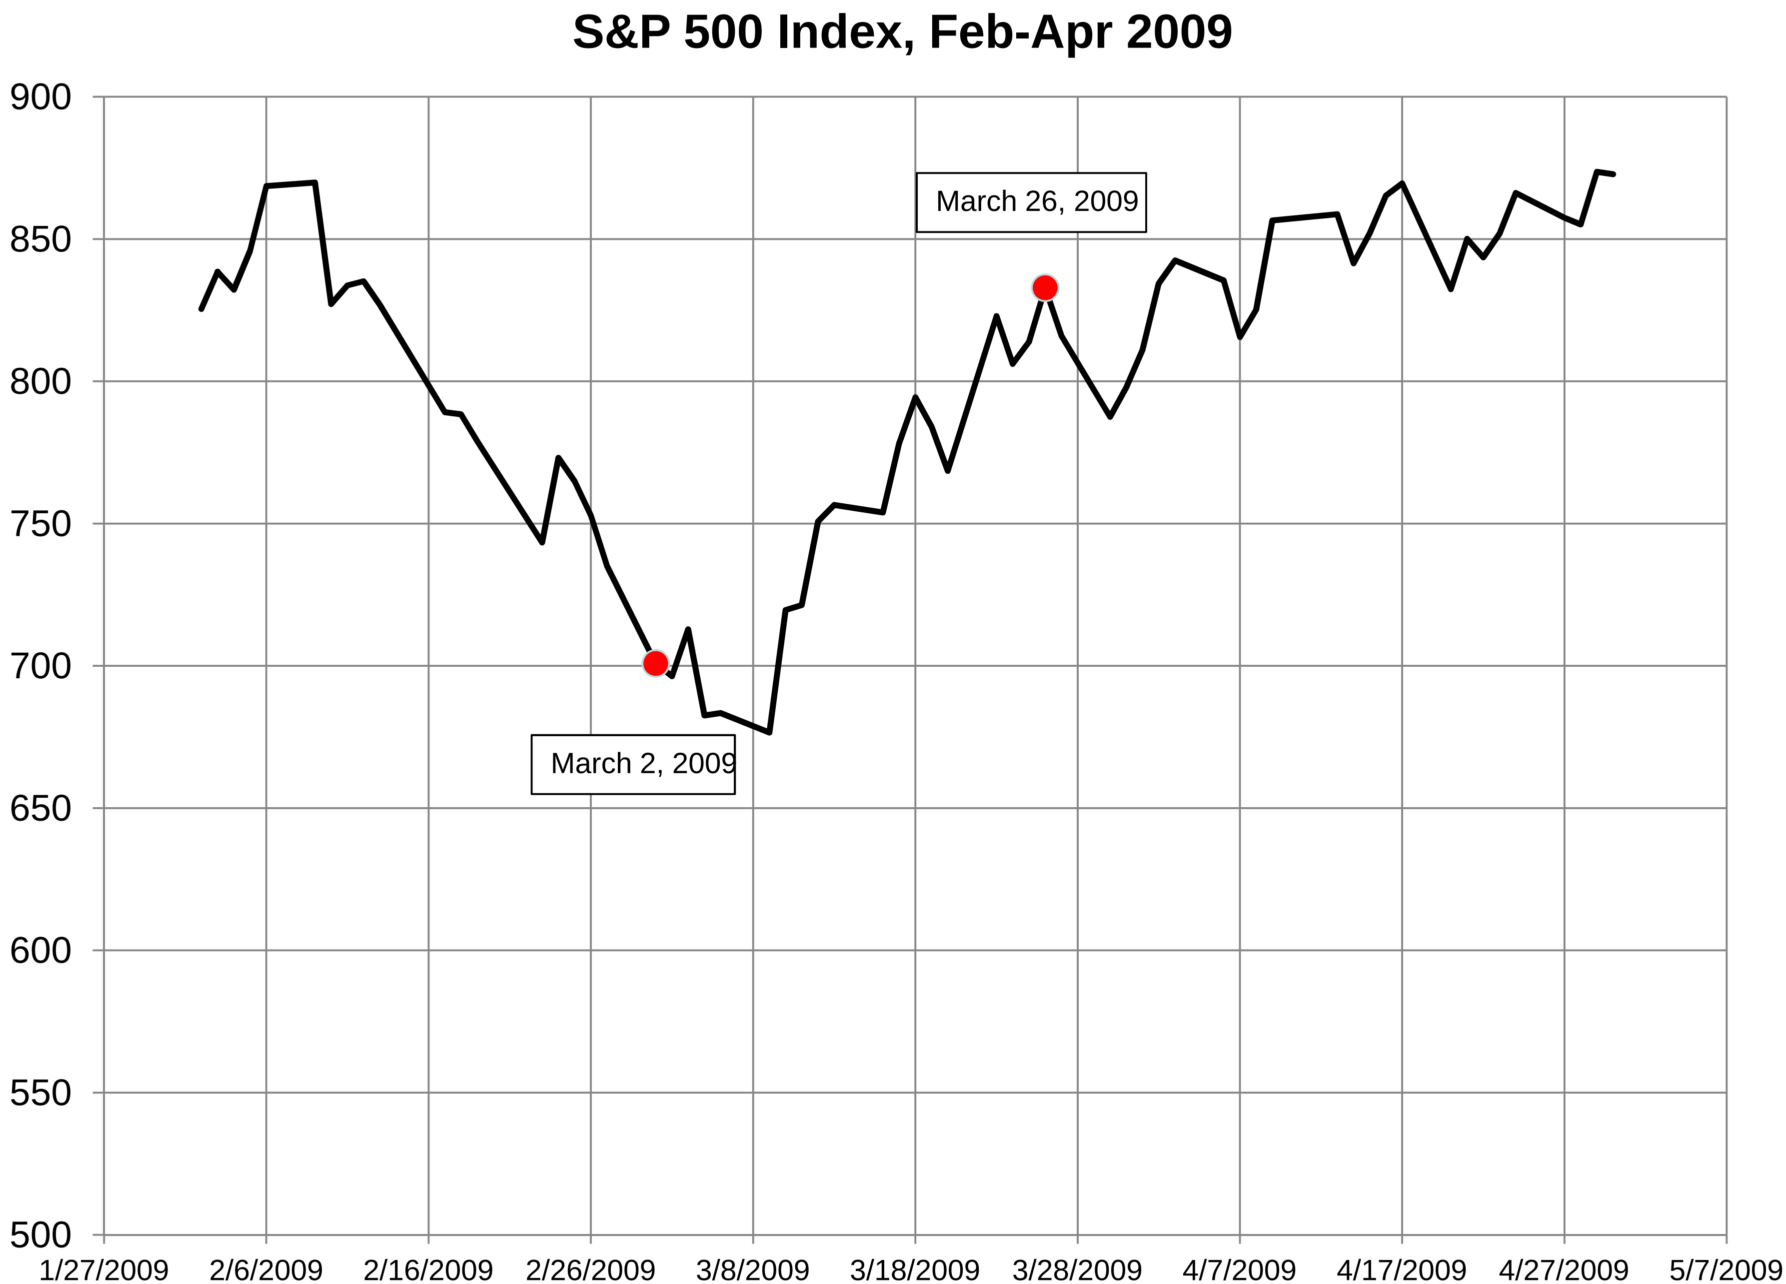
S&P 500 Index, Feb-Apr 2009
900
850
800
750
700
650
600
550
500
1/27/2009
2/6/2009
2/16/2009
2/26/2009
3/8/2009
3/18/2009
3/28/2009
4/7/2009
4/17/2009
4/27/2009
5/7/2009
March 26, 2009
March 2, 2009

Relative Demand Intensities, March 2, 2009

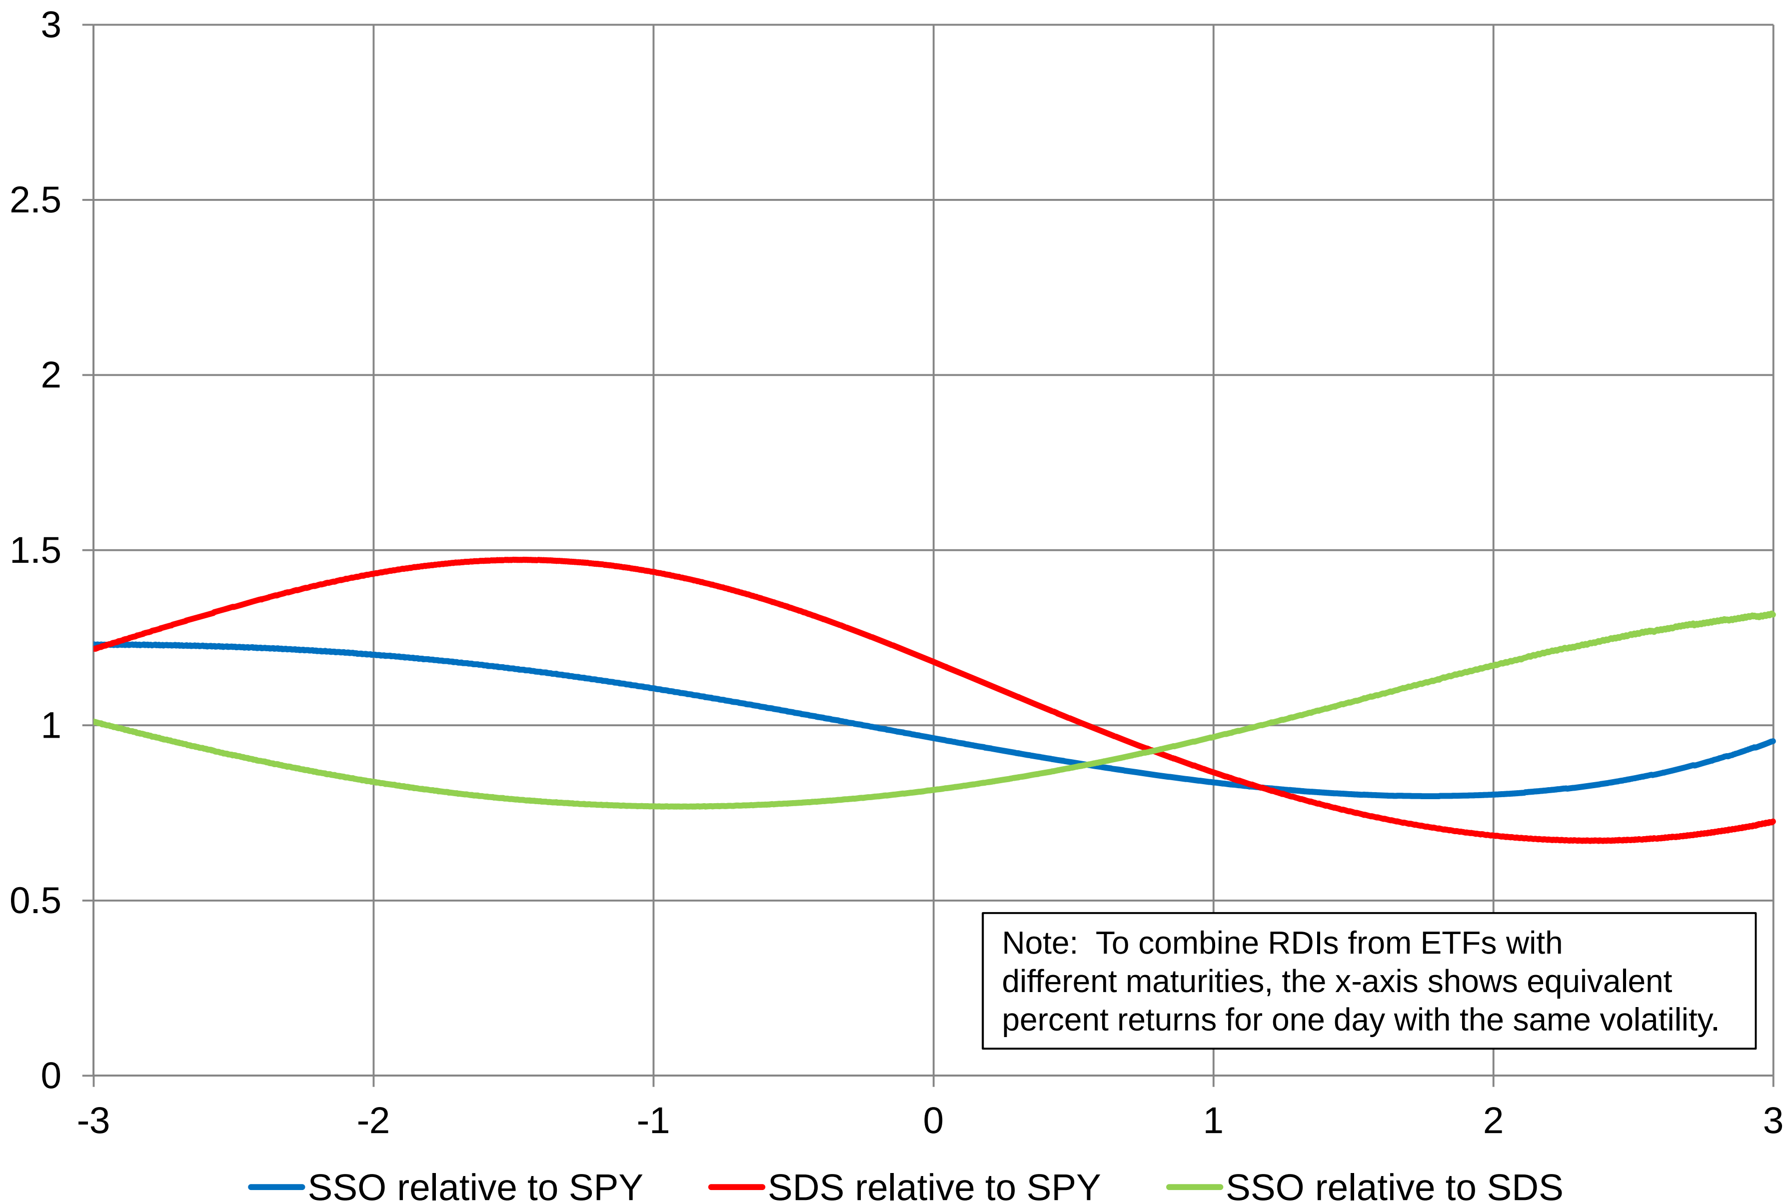

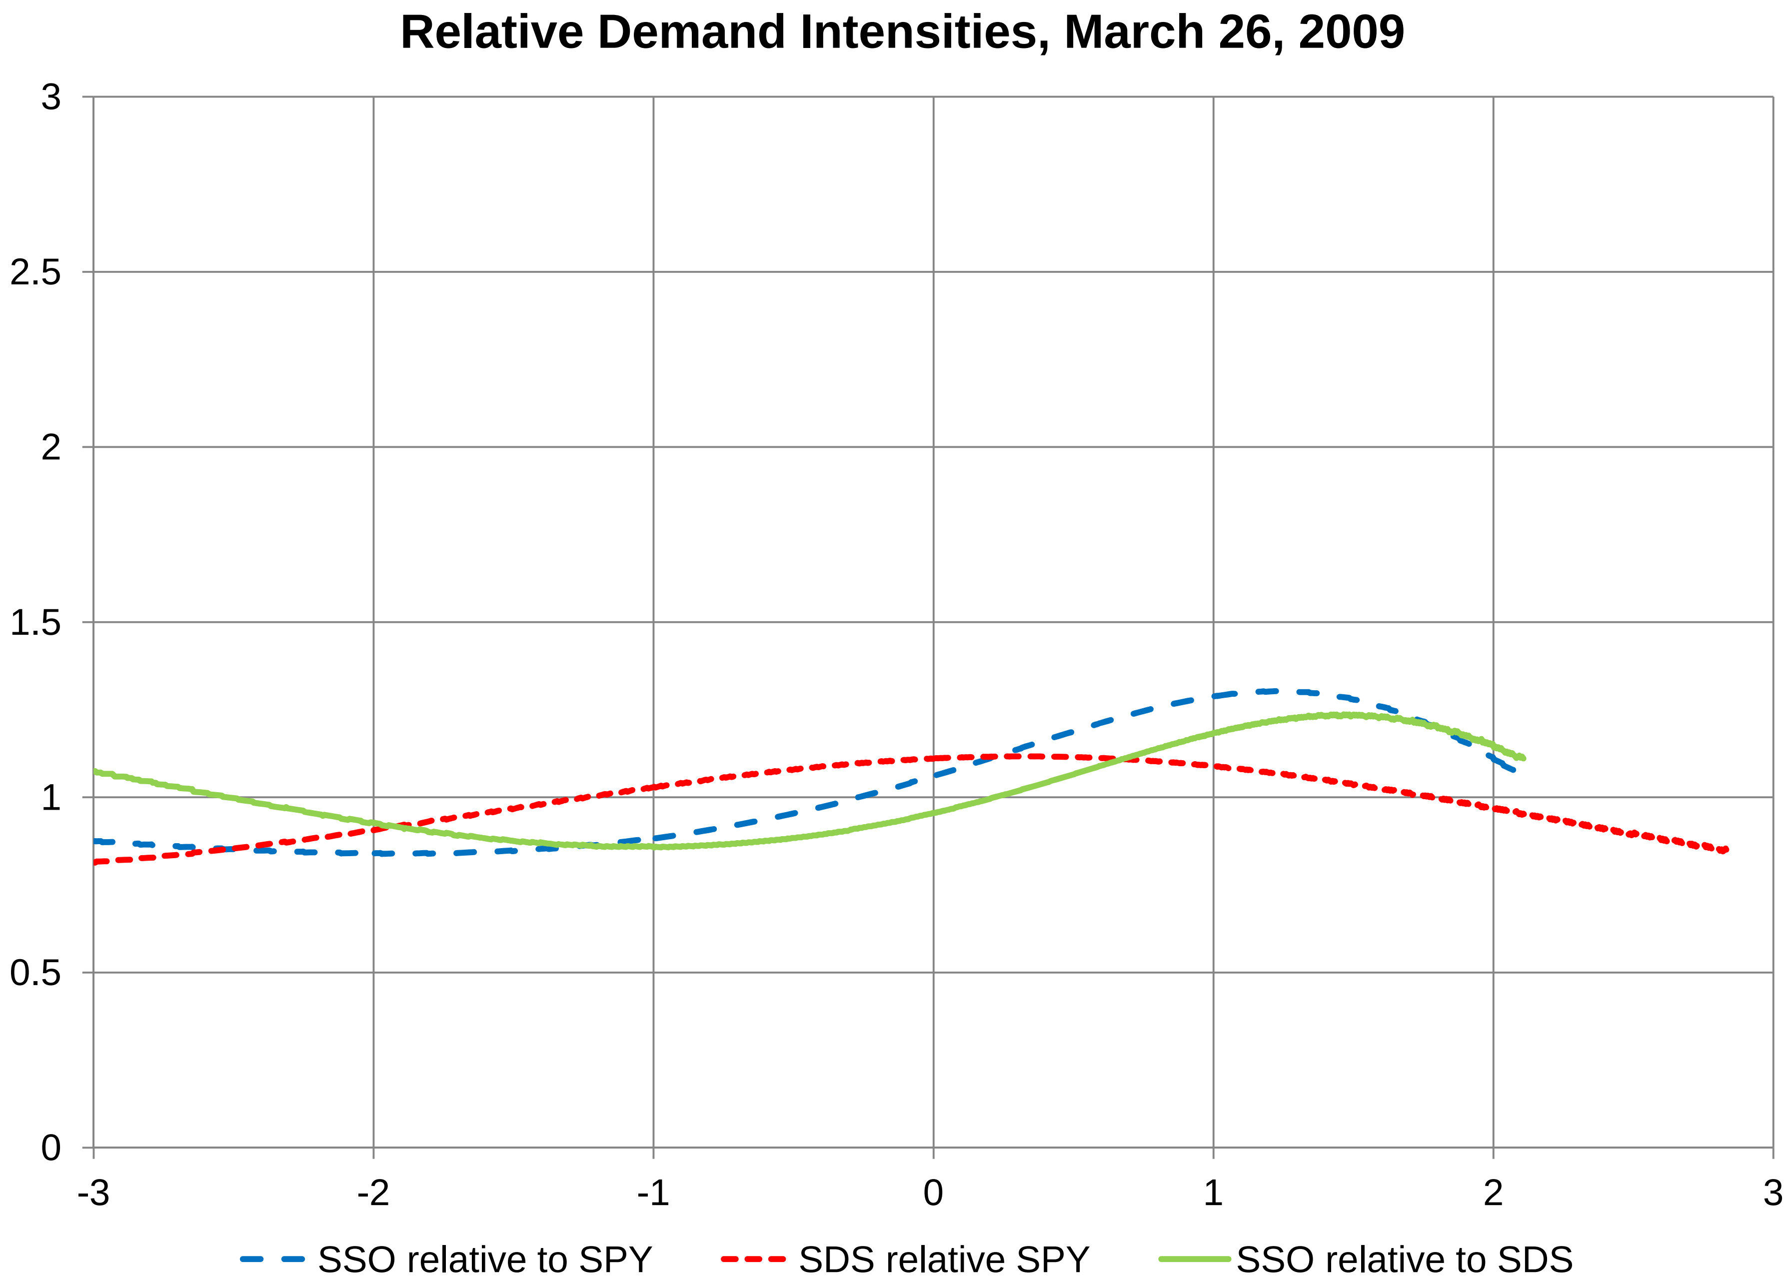

Relative Demand Intensities, March 26, 2009
3
2.5
2
1.5
1
0.5
0
-3
-2
-1
0
1
2
3
SSO relative to SPY
SDS relative SPY
SSO relative to SDS

A Trading Strategy?

The RND is invariably negatively skewed. Does less negative skewness indicate that investors are either feeling more confident or more bullish? If so, it might be a signal to go long.

Fahd has explored this question. He considers going long one of the ETFs on days when the (75-50-25) measure of RND is less negative than average and staying out of the market when the skewness measure is more negative than average.

	SPY	SSO	SDS	SPX
Annual return	19.1%	6.9%	20.1%	24.9%
Standard deviation	11.1%	12.6%	14.3%	22.4%
Return / Std dev	1.72	0.55	1.41	1.11
% of periods invested	56%	57%	58%	100%

A Trading Strategy?

Cumulative Returns:

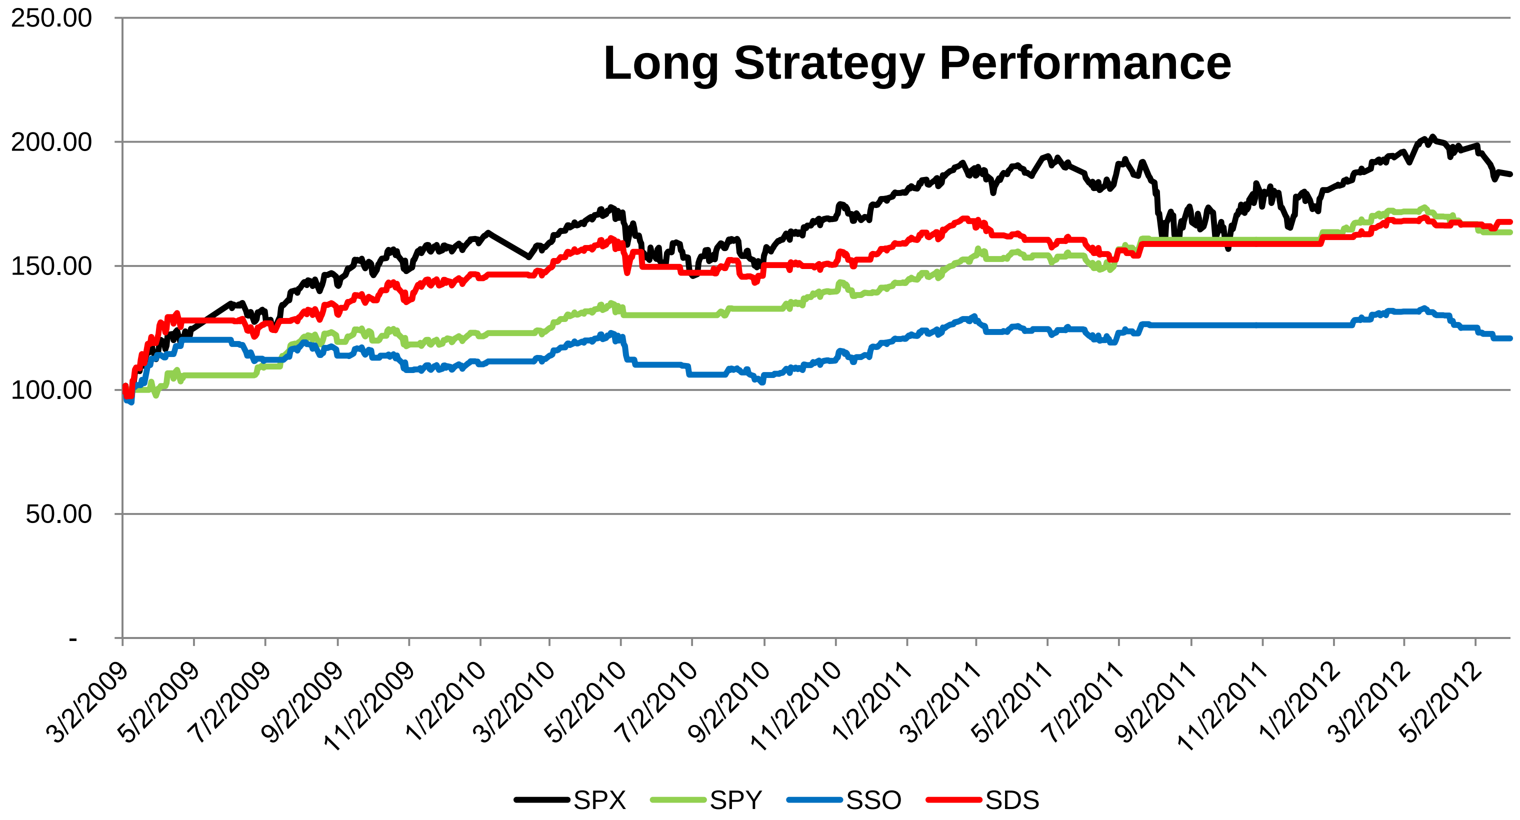

Theoretical Trading Strategy (untested!)

The Relative Demand Intensity curves show how much investors in one market value a particular future state of the world (stock market level) relative to investors in a different market, for example, buyers of the double long ETF SSO relative to those who buy the SDS double short ETF.

In theory, there is an arbitrage trade: Sell the contingent payoff in the market where it is priced high and buy it in the market where it is cheap.

On March 2, 2009 SDS investors valued a payoff in the case that the market goes down fairly sharply over the next 8 weeks much more than SSO investors did. The difference was even larger between SDS and SPY investors.

The trade: Buy exposure to that outcome with SPY or SSO options and sell it with SDS options. Use "butterfly spreads" in both markets to achieve those targeted exposures.

CAUTION:

"In theory there is not much difference between 'theory' and 'practice'.

But in practice there is."

Conclusions

The Risk Neutral Density contains a large amount of valuable information about market expectations and risk preferences.

We feel the solution to the "pricing kernel puzzle" is that investors have different price expectations and different risk tolerance, that "representative agent" models don't capture.

The first major challenge, how to extract a viable RND, has been well-explored and good methods are now in use.

The second major challenge, how to separate the market's true returns expectations from the risk premia, remains a challenge.

Our comparison of RNDs from Exchange Traded Funds with different exposures to the same underlying index gives some insights. We have suggested the ratio of the RNDs, which we call **Relative Demand Intensity**, as a way to measure how different investor clienteles value the same future outcome differently.

Insights can be gained

- from the shape of the RDIs over future market levels
- from the changes in RDIs in response to market moves and other factors

The "Bottom Line" Question

<u>Are there profitable trading strategies here?</u>

That is a good question to research!

(Please let me know what you find!)