The Recovery Theorem

September 20, 2013

NYU Stern – Federal Reserve Conference
NYU Stern School of Business

Steve Ross

Franco Modigliani Professor of Financial Economics
MIT

Managing Partner
Ross, Jeffrey & Antle LLC

The Distribution of Returns

- The generic approach to estimating return distributions in financial markets is purely statistical
- Typically we assume that the distribution of returns lies in some parametric class and use historical observations to estimate the unobserved parameters
- For example, we might assume that log stock returns follow are normally distributed and estimate the historical risk premium
- To estimate the standard deviation we use historical standard deviations or the current VIX
- To apply the results for financial decision making, e.g., asset allocation or any other policy decision, we assume that future returns are drawn from the historically estimated distribution
- But, the payoffs of financial instruments extend into the future and surely their prices are informative about the subjective distributions of the payoffs
- We extract such information in the fixed income markets

Predicting Interest Rates

- We use forward rates to tell us about anticipated future spot rates
- Forward rates aren't unbiased predictions of future spot rates because they have an embedded risk premium
- To separate the risk premium from the spot forecast, we usually estimate a parametric interest rate model
- Such models assume a distribution for interest rates, model the risk premium, then use historical data to estimate the relevant parameters
- This is sometime augmented by using the current yield curve to help fit the model

Forecasting the Equity Markets

- For equities we don’t even ask the market for a single spot forecast like we do with forward interest rates
- In the equity markets we:
  - Use historical market returns and the historical premium over risk-free returns to predict future returns
  - Build a model (e.g., a dividend/yield model) to predict stock returns
  - Use the martingale measure from options or some ad hoc adjustment to it as though it was the same as the natural probability distribution
  - Survey market participants and institutional peers
- What we want to do in the equity markets – and the fixed income markets – is find the market’s subjective distribution of future returns
- At the least this would open up a host of different empirical possibilities
- Let’s look at the derivatives market and option pricing models

The Binomial Model

- Arguably, the Black-Scholes-Merton model and its close cousin, the Binomial model, are the most successful models in economics (and, perhaps, in all of the social sciences):

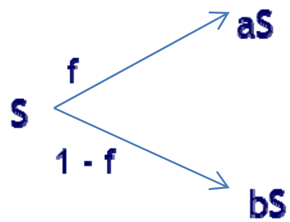

$$b < 1 + r < a$$

where f is the natural jump probability and r is the risk free interest rate

- The absence of arbitrage implies the existence of positive (Arrow Debreu (AD)) prices for pure contingent claims, i.e., digitals that pay $1 in state a or b:

$$p(a) = \left(\frac{1}{1+r}\right)\pi \quad and \quad p(b) = \left(\frac{1}{1+r}\right)(1-\pi)$$

where

$$\pi = \frac{(1+r)-b}{a-b} \quad and \quad 1-\pi = \frac{a-(1+r)}{a-b}$$

The Binomial Model

- Notice that the sum of the prices is the discount factor (1/(1+r)), and that

$$p(a)a + p(b)b = 1$$

- In this formulation π is called the risk neutral probability and prices are the discounted expected value of payoffs using the risk neutral probabilities

- Equivalently, no arbitrage and spanning imply that there is a unique portfolio of the stock and the bond that replicates – delta hedges - the payoff of the option and the cost of this portfolio is simply the price of the derivative

- Strikingly, the natural probability, f, plays no role whatsoever in this theory for pricing derivatives

- Consequently, there is no way to use derivatives prices to find any information about the underlying natural distribution, f

Black-Scholes-Merton

- The Binomial model converges to a lognormal diffusion as the step size gets smaller

$$\frac{dS}{S} = \mu dt + \sigma dz$$

- The Binomial formula for the price of an option also converges to the famous Black-Scholes-Merton risk neutral differential equation

$$\frac{1}{2}\sigma^2 S^2 \frac{\partial^2 P}{\partial S^2} + rS\frac{\partial P}{\partial S} - rP + \frac{\partial P}{\partial t} = 0$$

yielding the equally famous Black-Scholes-Merton solution for a call option

$$P(S,t) = SN(d_1) - Ke^{-r(T-t)}N(d_2)$$

where N(∘) is the cumulative normal, K is the strike, and

$$d_1 = \frac{\ln\left(\frac{S}{K}\right) + \left(r + \left(\frac{1}{2}\right)\sigma^2\right)(T-t)}{\sigma\sqrt{T-t}} \quad and \quad d_2 = d_1 - \sigma\sqrt{T-t}$$

Black-Scholes-Merton

- As expected, the Black-Scholes-Merton solution is also the discounted expected value of the payoff on a call using the risk neutral probability distribution, i.e., under the assumption that the stock drift is the risk free rate
- Since the expected return on the stock, μ, doesn't enter the formula, option prices are independent of that component of the natural distribution of the stock
- This is both a blessing and a curse;
- It means that we can price derivatives without knowing the very difficult to measure parameters of the underlying process
- But, conversely, with these pricing models, the market for derivatives is of no use to us for determining the natural distribution
- The challenge, then, is to build a different class of model that will have the ability to link option prices to the underlying parameters of the asset process while at the same time being consistent with risk neutral pricing

The Options Market

- A put option is insurance against a market decline with a deductible; if the market drops by more than the strike, then the put option will pay the excess of the decline over the strike, i.e., the strike acts like a deductible
- (A call option on the market is a security with a specified strike price and maturity, say one year, that pays the difference between the market and the strike iff the market is above the strike in one year)
- The markets use the Black-Scholes-Merton formula to quote option prices, and volatility is an input into that formula
- The implied volatility is the volatility that the stock must have to reconcile the Black-Scholes-Merton formula price with the market price
- Notice that the market isn't necessarily using the Black-Scholes-Merton formula to price options, only to quote their prices

The Volatility Surface

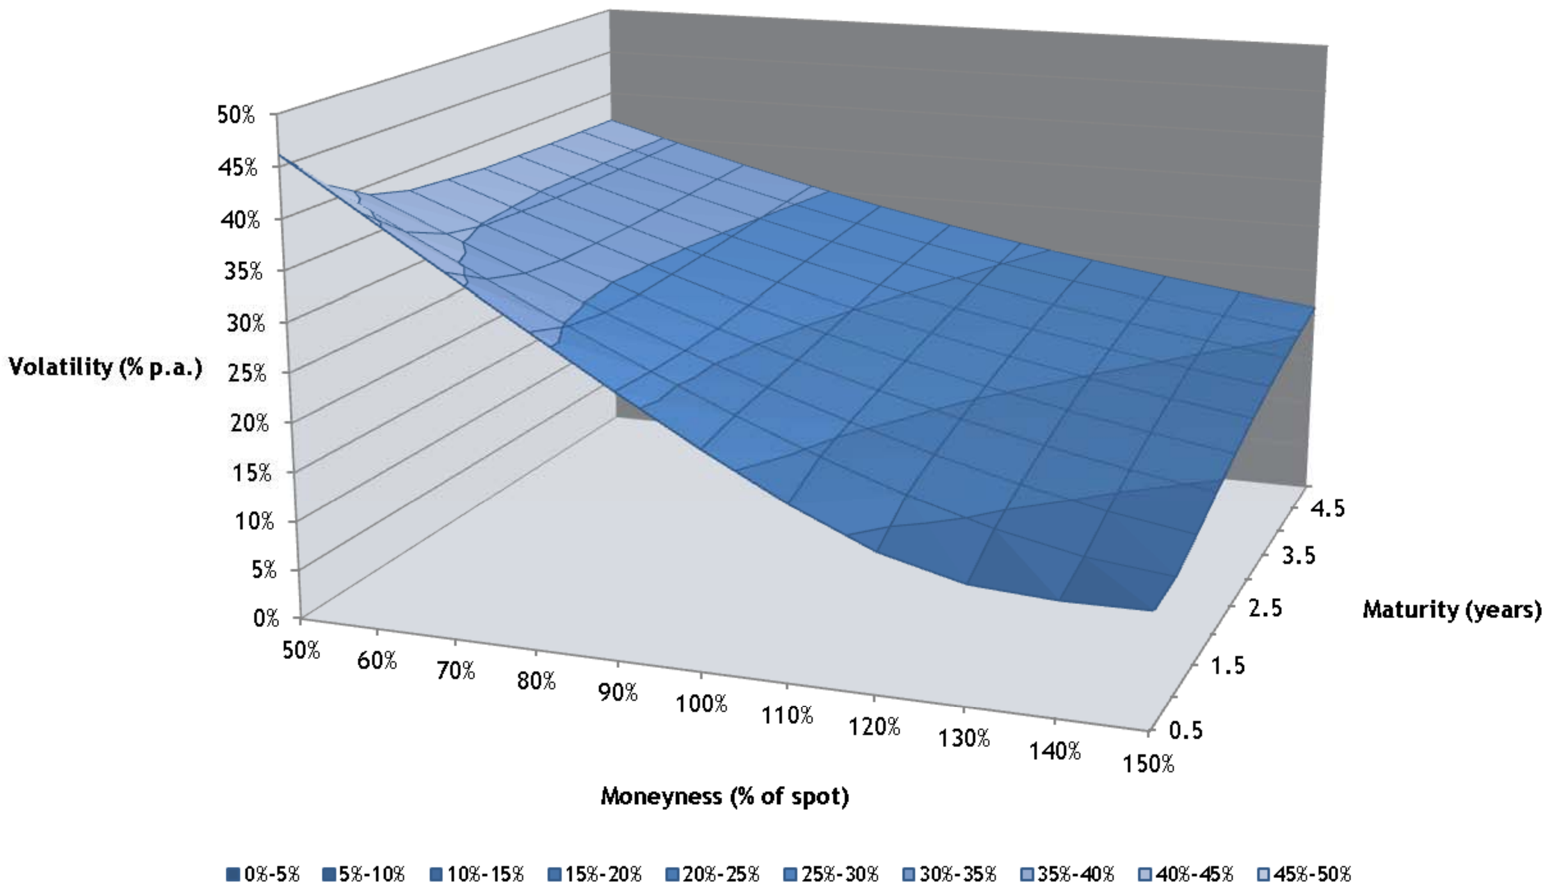

Surface date: January 6, 2012

Implied Vols, Risk Aversion, and Probabilities

- Like any insurance, put prices are a product of three effects:

  Put price = Discount Rate x Risk Aversion x Probability of a Crash

- But which is it – how much of a high price comes from high risk aversion and how much from a high chance of a crash?

- The Binomial model and Black-Scholes-Merton not only provide no answer to this question, they also suggest that no answer is possible and, to some extent, we've lost the intuitions of insurance in derivative theory

- In fact, much is made of the distinction between risk neutral probabilities which are inferred from option prices, and the latent natural probabilities, and generally research carefully and rightfully separates these different concepts

- Nonetheless, one often encounters a careless blurring of this distinction

- For example, risk aversion is sometimes ignored and the skew of the vol surface is casually interpreted as proof that tail probabilities are large

Contingent Forward Prices

- To recover the true distribution of stock returns we use a new concept, contingent forward prices, i.e., the forward prices that would prevail from any value of the market and not just the current value
- Implicit in option prices are the prevailing forward prices for digital securities that pay off if the market is in a given range, e.g., between 1700 and 1750 in one year
- Suppose the market is in some state i (for simplicity we will index states by the stock price alone) and we want to know the price of a digital security that pays $1 if the market is in state j in a year, i.e., the contingent forward price, p_{ij}
- For example, state i could be S&P 500 = 1500 and state j could be S&P 500 = 1550
- Like forward rates in the fixed income market, contingent forward prices are prices that we can 'lock in' today for buying a digital paying, say, if the market is at 1700 in one year and drops to 1400 in the next six months
- First, though, we have to find these contingent forward prices and since the current state is c = 1600, rather than some arbitrary state i and since we don't have forward markets we have to demonstrate that this is possible

Contingent Forward Prices

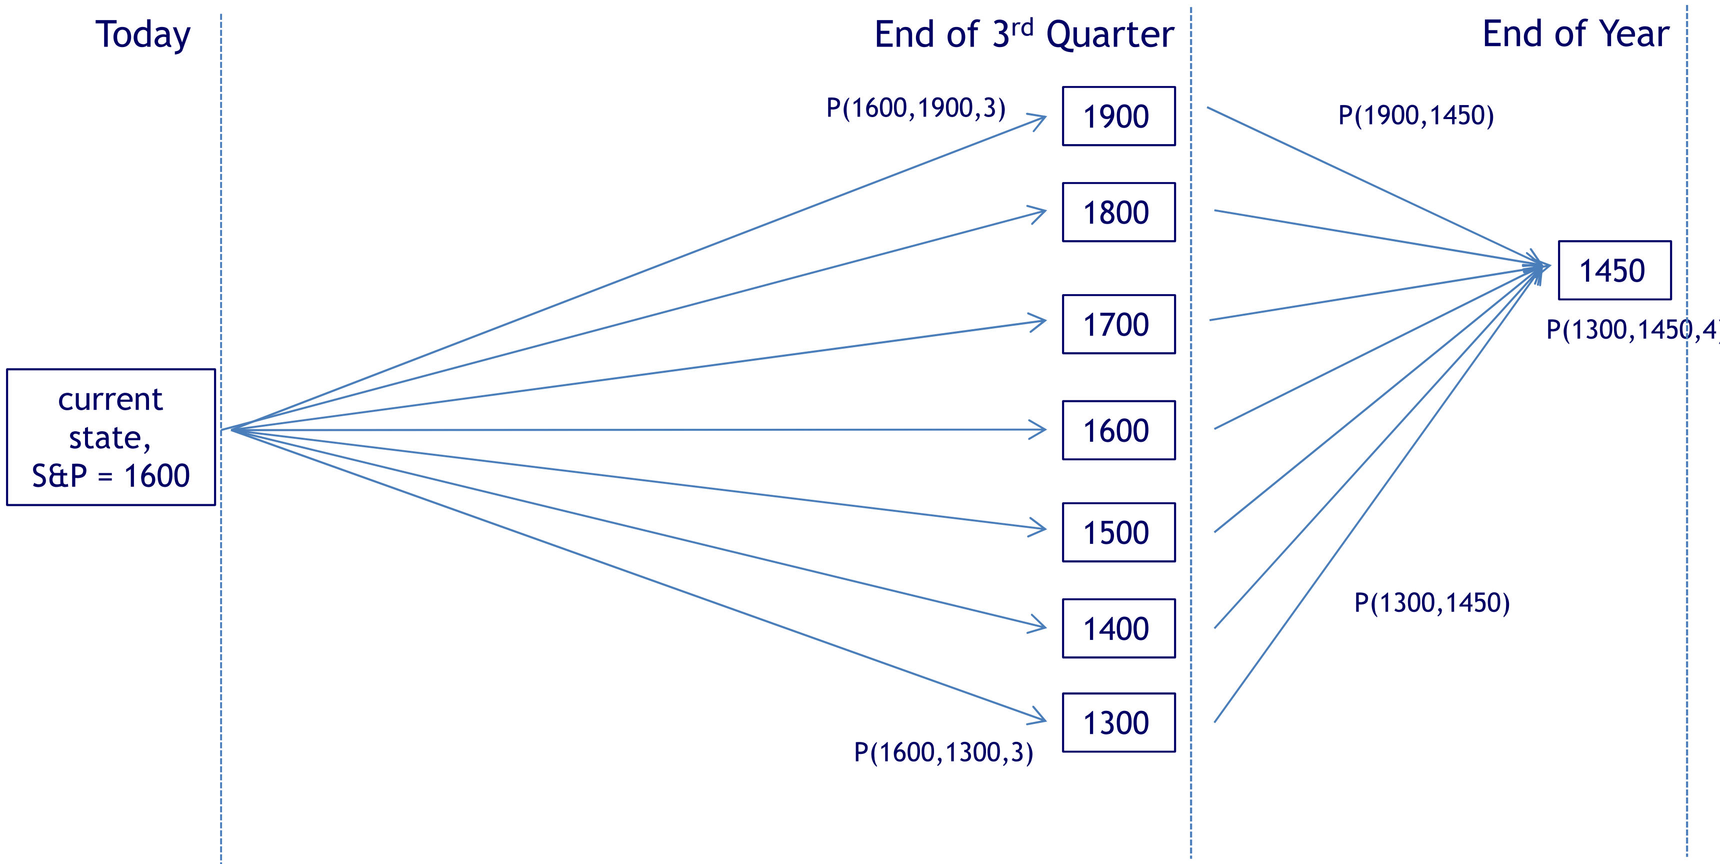

The Recovery Theorem

- A contingent forward security is a form of insurance, hence:

$$\text{Contingent Forward Price} = p_{ij} = \delta \varphi_{ij} f_{ij}$$

where δ is the market's average discount rate, ϕ(i,j) is the pricing kernel and captures risk aversion, and f_{ij} is the natural probability of transiting from state i to state j

- Knowing $P = [p_{ij}]$ there is no way to disentangle the kernel from the probabilities
  - There are $2m^2 + 1$ unknowns and only $m^2 + m$ equations
- Suppose, though, that the kernel is transition independent, i.e., that it has the form φ(j)/φ(i)
- For <u>example</u>, for a representative agent with additively separable preferences

$$\varphi_{ij} = \delta \frac{U'(C(j))}{U'(C(i))} = \delta \frac{\varphi(j)}{\varphi(i)}$$

- Now we can separate φ from f

The Recovery Theorem

- Letting D be the diagonal matrix with the m kernel values, $\phi(j)$, on the diagonal, the pricing equation can be written as:

$$P = \delta D^{-1} F D$$

- Since the probabilities in each row have to add up to one (m additional equations), if e is the vector of ones, then we have:

$$Fe = e$$

- Rearranging the pricing equation we have:

$$F = (1/\delta)\, DPD^{-1}$$

and since $Fe = e$,

$$DPD^{-1}e = \delta e,$$

or

$$Px = \delta x,$$

where $x = D^{-1}e$ is a vector whose elements are the inverses of the kernel

The Recovery Theorem

- From the Perron-Frobenius Theorem (if P is irreducible), we know this equation has a unique positive eigenvector solution, x, and an associated positive eigenvalue δ
- From x = (1/φ(i)) we have the kernel and, given δ, we can now solve for the natural probabilities, F

$$f_{ij} = \left(\frac{1}{\delta}\right)\left(\frac{\varphi(i)}{\varphi(j)}\right)p_{ij}$$

Recovery and the Binomial and Black-Scholes-Merton Models

- There are three key features that permitted us to recover the natural distribution:
- First, we assumed the kernel is transition independent – this reduces the kernel to m unknowns
- Second, we introduced state dependence – this gives m^2 equations, not m
- Third, the state space was finite
- For the Black-Scholes-Merton model the state space is unbounded and it can be shown that the there is a one parameter family of solutions which cannot be distinguished, i.e., we can recover

$$\mu - \varphi\sigma^2$$

but not μ and φ separately

Some Extensions

- If we attempt to apply the Recovery Theorem to a Black-Scholes-Merton environment we discover that there are multiple solutions; any exponential is an eigenvector of the state price transition function
- There are two sufficient conditions that permit extensions to continuous spaces
- First, if the state space is bounded we can apply the Krein-Rutman Theorem which is a direct extension of Perron-Frobenius to continuous states if the space has an interior (e.g., closed, convex, and with a boundary and the transition function is a compact linear operator and strongly positive (i.e., $Pv \gg 0$ if $v > 0$)
- Compactness may be a good assumption for interest rate processes
- Second, if the kernel is bounded above and below we can prove that there is a unique positive eigenvalue with an associated positive eigenfunction, φ

$$0 < a \leq \varphi \leq b < \infty$$

Applying the Recovery Theorem – A Three Step Procedure

- Pick a date
- Step 1: Use option prices to determine the forward prices
- Step 2: Use the forward prices to determine the contingent forward prices
- Step 3: Apply the Recovery Theorem to determine the risk aversion and the market's natural probability distribution of equity returns

Step 1: Forward (Digital) Prices

- From options of different maturities and strikes we can create and price the digital securities that pay off if in one year the market is between, say, 1550 and 1560
- A one year bull butterfly spread, composed of one call with a strike of 1550, one with a strike of 1560, and selling two calls with strikes of 1555 pays off iff the market is between 1550 and 1560
- Notice that the prices are peaked near the current value for the market and fall off and spread at longer maturities
- Notice, too, that the prices are lowest for large moves, higher for big down moves than for big up moves, and that after an initial fall prices tend to rise with tenor

Forward (Digital) Prices

Pure Security Prices

Market Scenario	Tenor 0.25	0.50	0.75	1.00	1.25	1.50	1.75	2.00	2.25	2.50	2.75	3.00
-35%	$0.005	$0.023	$0.038	$0.050	$0.058	$0.064	$0.068	$0.071	$0.073	$0.075	$0.076	$0.076
-29%	$0.007	$0.019	$0.026	$0.030	$0.032	$0.034	$0.034	$0.035	$0.035	$0.035	$0.034	$0.034
-23%	$0.018	$0.041	$0.046	$0.050	$0.051	$0.052	$0.051	$0.050	$0.050	$0.049	$0.048	$0.046
-16%	$0.045	$0.064	$0.073	$0.073	$0.072	$0.070	$0.068	$0.066	$0.064	$0.061	$0.058	$0.056
-8%	$0.164	$0.156	$0.142	$0.128	$0.118	$0.109	$0.102	$0.096	$0.091	$0.085	$0.081	$0.076
0%	$0.478	$0.302	$0.234	$0.198	$0.173	$0.155	$0.141	$0.129	$0.120	$0.111	$0.103	$0.096
9%	$0.276	$0.316	$0.278	$0.245	$0.219	$0.198	$0.180	$0.164	$0.151	$0.140	$0.130	$0.120
19%	$0.007	$0.070	$0.129	$0.155	$0.166	$0.167	$0.164	$0.158	$0.152	$0.145	$0.137	$0.130
30%	$0.000	$0.002	$0.016	$0.036	$0.055	$0.072	$0.085	$0.094	$0.100	$0.103	$0.105	$0.105
41%	$0.000	$0.000	$0.001	$0.004	$0.009	$0.017	$0.026	$0.036	$0.045	$0.053	$0.061	$0.067
54%	$0.000	$0.000	$0.000	$0.000	$0.000	$0.000	$0.001	$0.001	$0.002	$0.002	$0.003	$0.003

Priced using the SPX volatility surface from April 27, 2011

- A pure security price is the price of a security that pays one dollar in a given market scenario for a given tenor - for example, a security that pays $1 if the market is unchanged (0% scenario) in 6 months costs $0.302

The Forward Price Surface

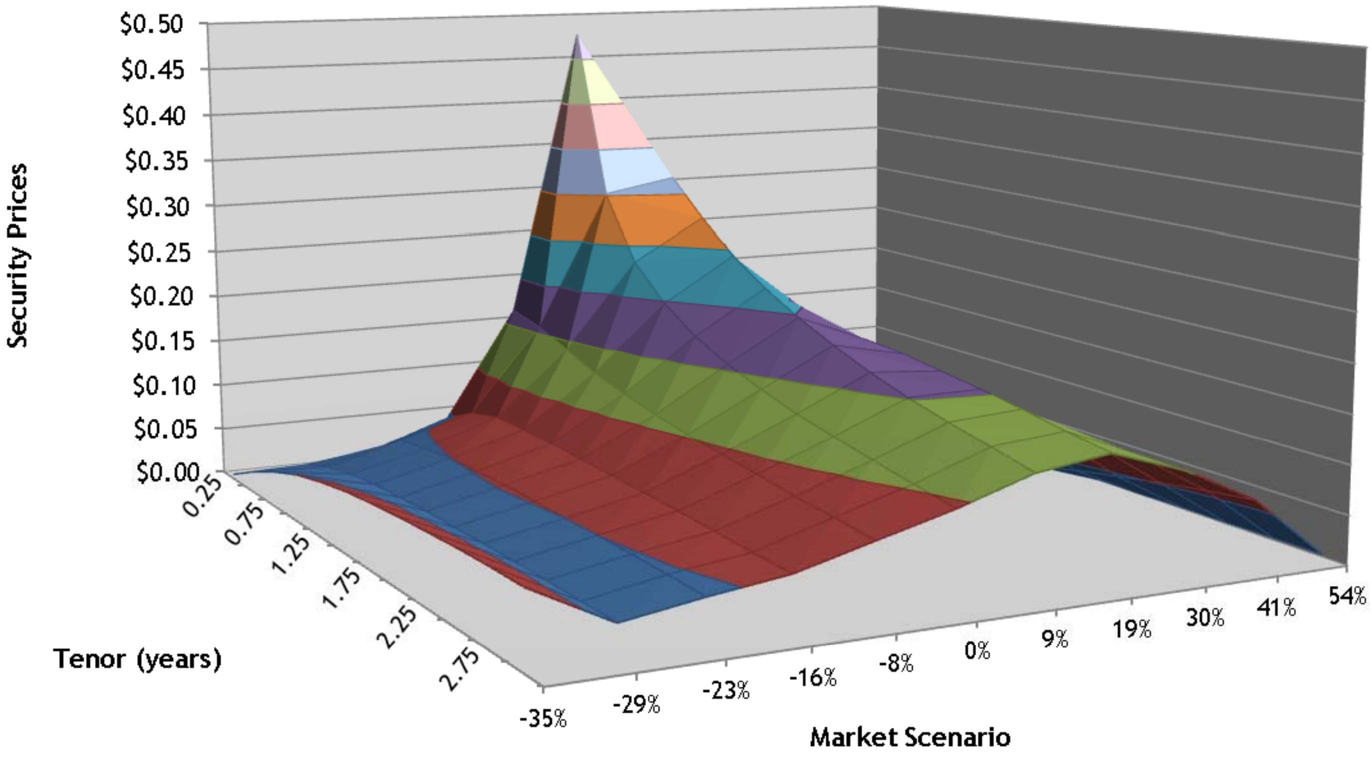

Priced using the SPX volatility surface from April 27, 2011

Step 2: Contingent Forward Prices

- We can now use the technique described earlier to determine the contingent forward prices, p_{ij}, from the forward prices

- The table on the following slide displays the resulting contingent forward prices

Contingent Forward Prices

Quarterly

Contingent Forward Prices

		Market Scenario Final Period										
		-35%	-29%	-23%	-16%	-8%	0%	9%	19%	30%	41%	54%
	-35%	$0.671	$0.241	$0.053	$0.005	$0.001	$0.001	$0.001	$0.001	$0.001	$0.000	$0.000
	-29%	$0.280	$0.396	$0.245	$0.054	$0.004	$0.000	$0.000	$0.000	$0.000	$0.000	$0.000
	-23%	$0.049	$0.224	$0.394	$0.248	$0.056	$0.004	$0.000	$0.000	$0.000	$0.000	$0.000
Market	-16%	$0.006	$0.044	$0.218	$0.390	$0.250	$0.057	$0.003	$0.000	$0.000	$0.000	$0.000
Scenario	-8%	$0.006	$0.007	$0.041	$0.211	$0.385	$0.249	$0.054	$0.002	$0.000	$0.000	$0.000
Initial	0%	$0.005	$0.007	$0.018	$0.045	$0.164	$0.478	$0.276	$0.007	$0.000	$0.000	$0.000
Period	9%	$0.001	$0.001	$0.001	$0.004	$0.040	$0.204	$0.382	$0.251	$0.058	$0.005	$0.000
	19%	$0.001	$0.001	$0.001	$0.002	$0.006	$0.042	$0.204	$0.373	$0.243	$0.055	$0.004
	30%	$0.002	$0.001	$0.001	$0.002	$0.003	$0.006	$0.041	$0.195	$0.361	$0.232	$0.057
	41%	$0.001	$0.000	$0.000	$0.001	$0.001	$0.001	$0.003	$0.035	$0.187	$0.347	$0.313
	54%	$0.000	$0.000	$0.000	$0.000	$0.000	$0.000	$0.000	$0.000	$0.032	$0.181	$0.875

Priced using the SPX volatility surface from April 27, 2011

- Contingent forward prices are the prices of securities that pay one dollar in a given future market scenario, given the current market range - for example, a security that pays $1 if the market is down 16% next quarter given that the market is down 8% in the current quarter costs $0.211

Step 3: Apply the Recovery Theorem

- Apply the Recovery Theorem and solve for the unknown probabilities and the risk aversions

- The next slides compare these predictions for a given date with the distribution obtained by bootstrapping monthly historical data

Risk Neutral Density
May 1, 2009

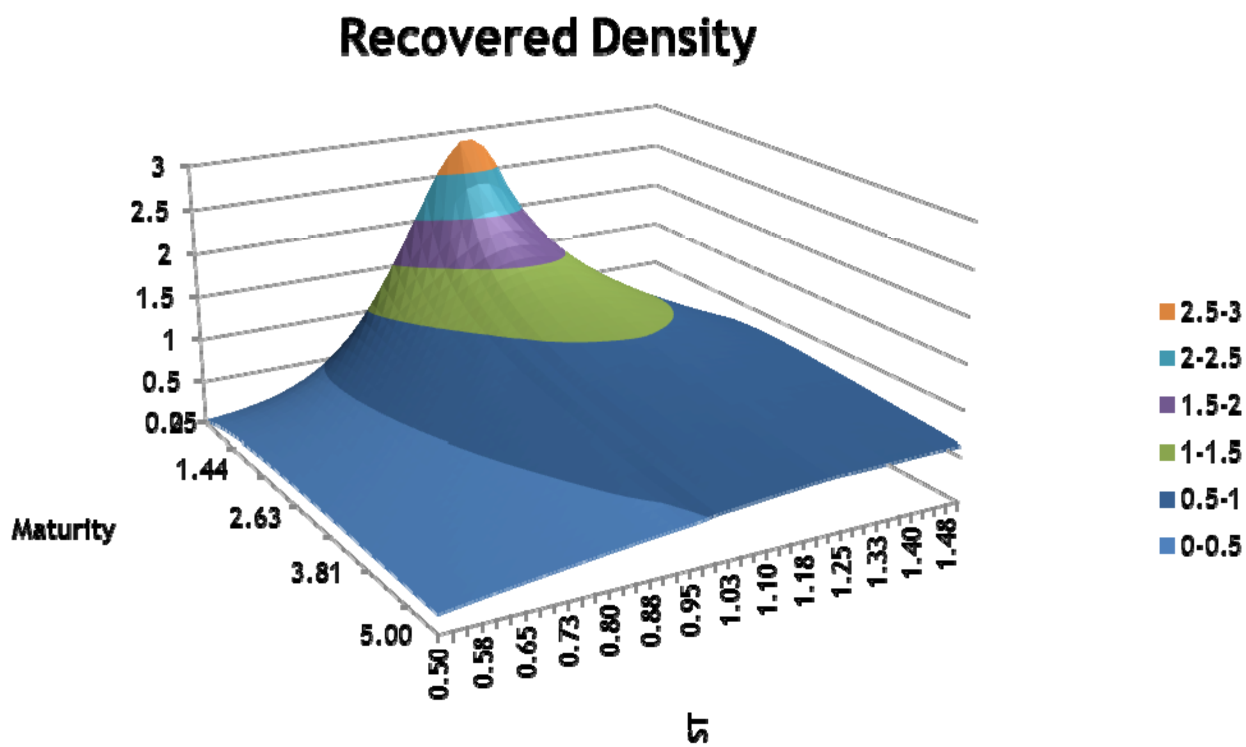

Implied Market Risk Aversion
May 1, 2009

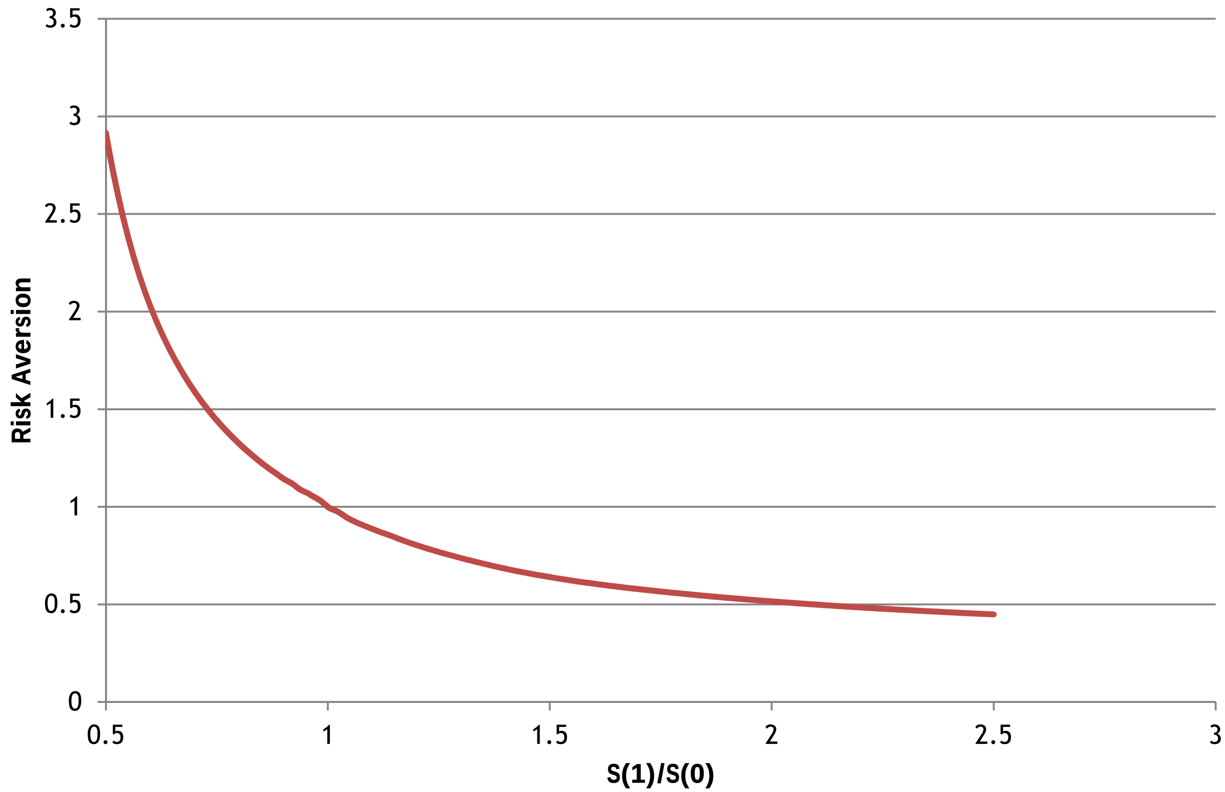

Recovered Density
May 1, 2009

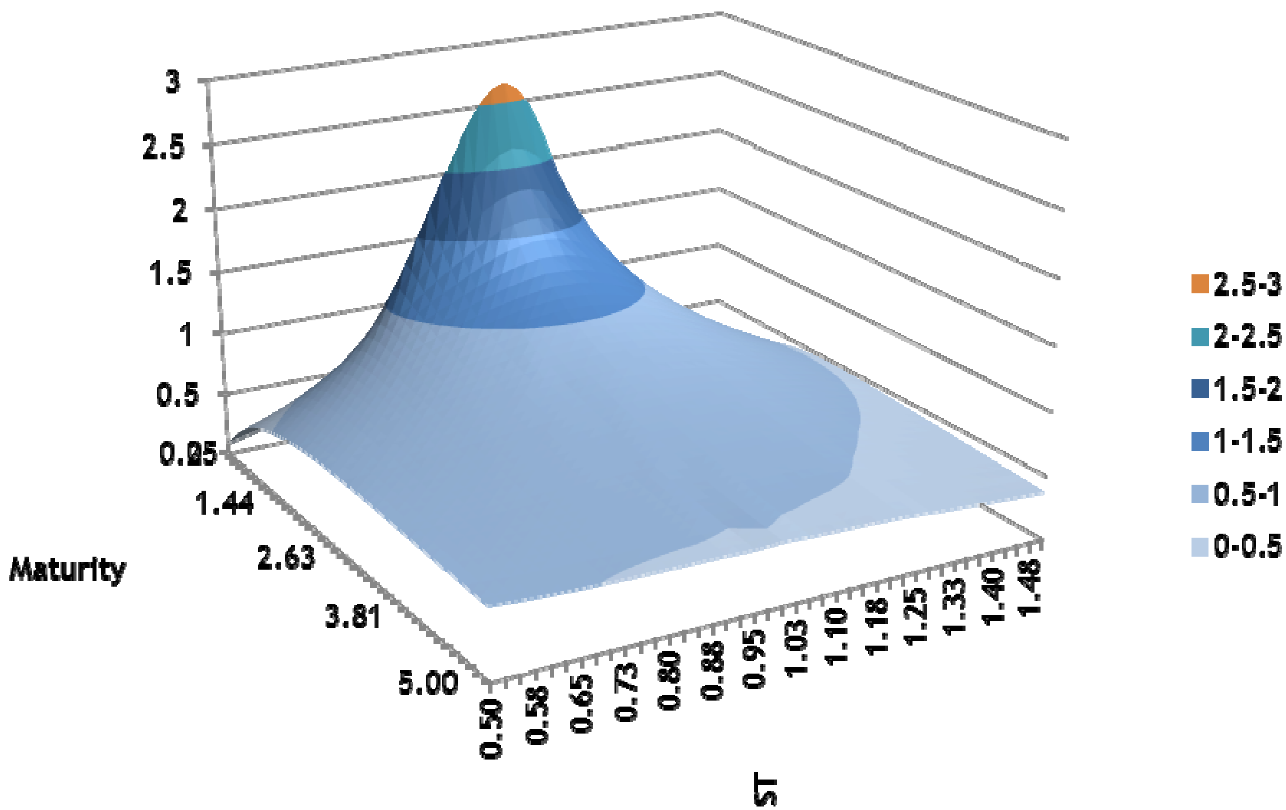

Recovered Probabilities vs. Risk-Neutral Probabilities
May 1, 2009

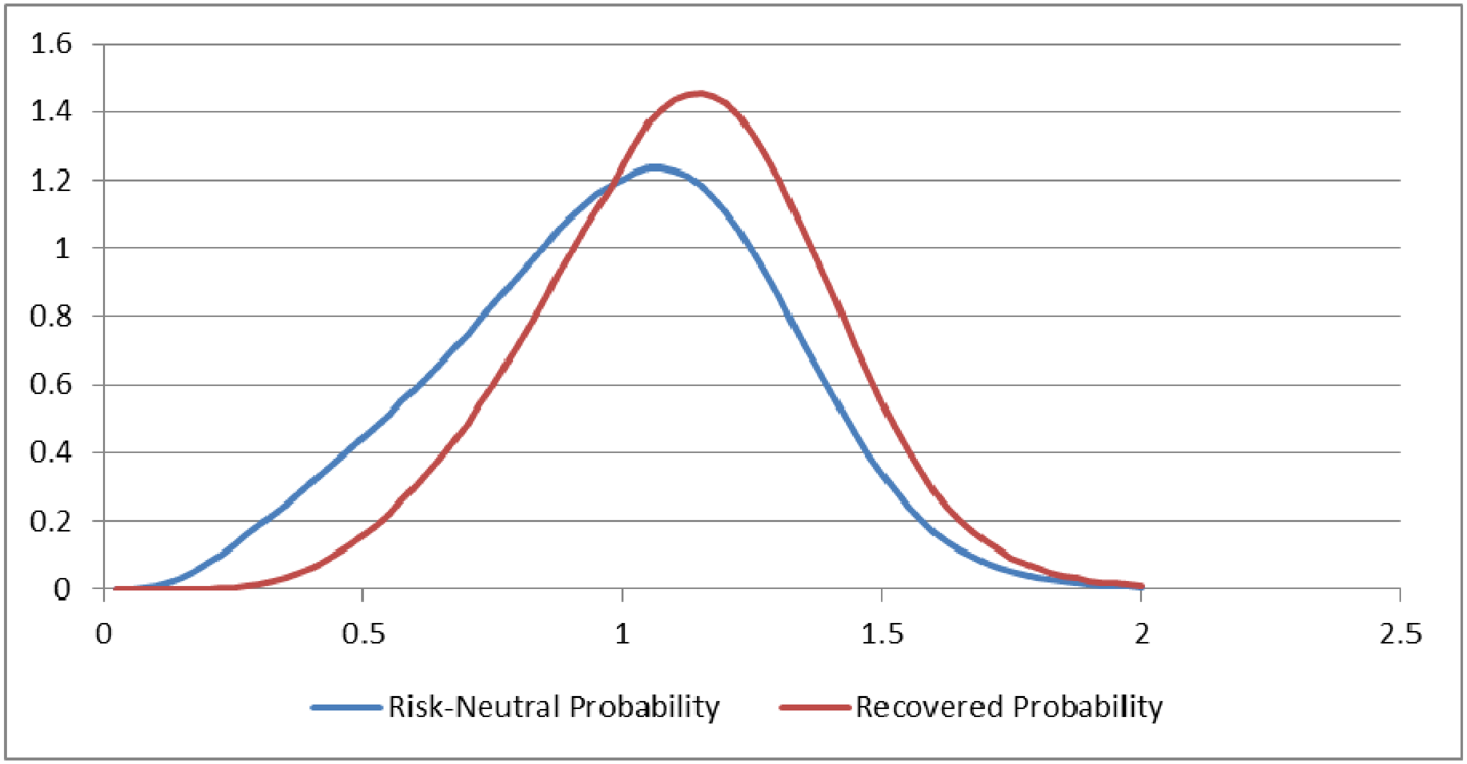

Recovered Probabilities vs. Historical Probabilities

May 1, 2009

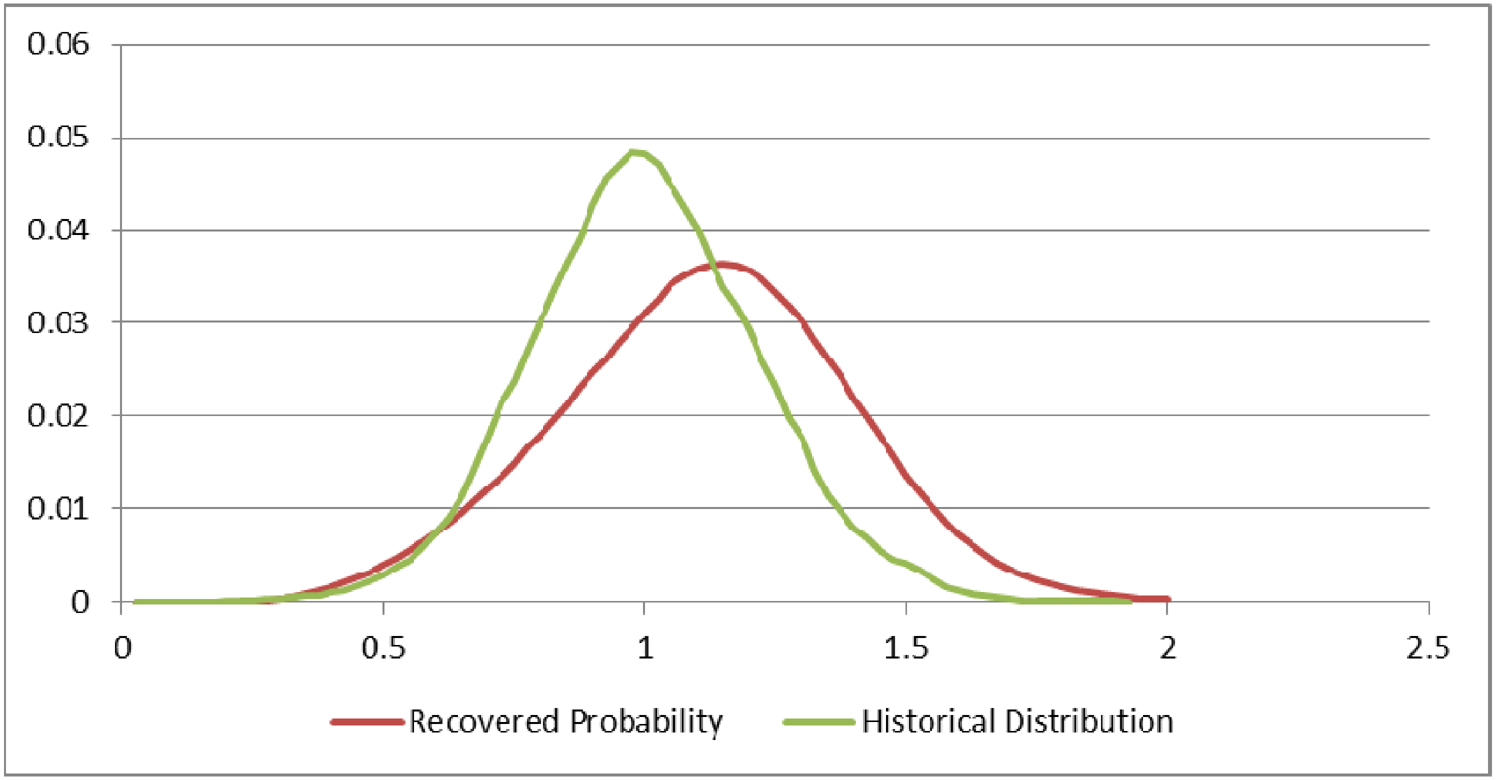

Recovered Cumulative vs. Historical Cumulative

May 1, 2009

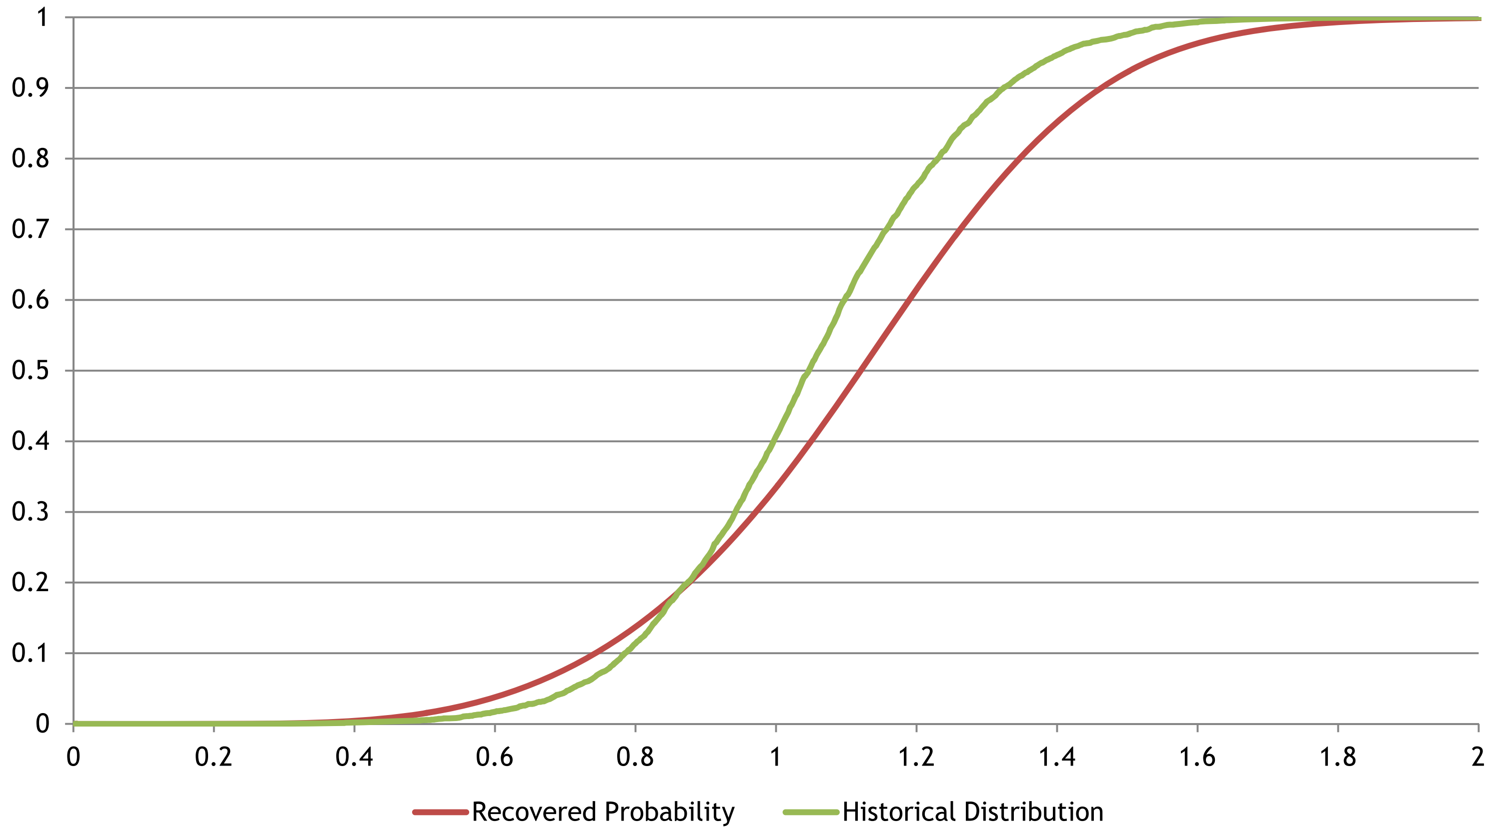

The Bootstrapped (Historical) and Recovered Probabilities

On May 1, 2009 for May 1, 2010

Market Scenario	Historical Bootstrap	Recovered Probabilities
-50%	0.005	0.017
-40%	0.017	0.041
-30%	0.045	0.083
-20%	0.114	0.146
-10%	0.233	0.234
-5%	0.315	0.288
0%	0.407	0.349
5%	0.507	0.416
10%	0.605	0.487
20%	0.762	0.631
30%	0.881	0.761
40%	0.947	0.861
50%	0.976	0.929

Recovered (May 1, 2009) and Historical Statistics

	Historical	Recovered
mean	8.40%	9.81%
excess return	3.06%	7.13%
sigma	15.30%	16.28%
Sharpe	0.20	0.44
rf	5.34%	1.08%
divyld	3.25%	2.67%
ATM ivol		32.84%

Actual Volatility vs. VIX
(May 1, 2009 – May 1, 2013)

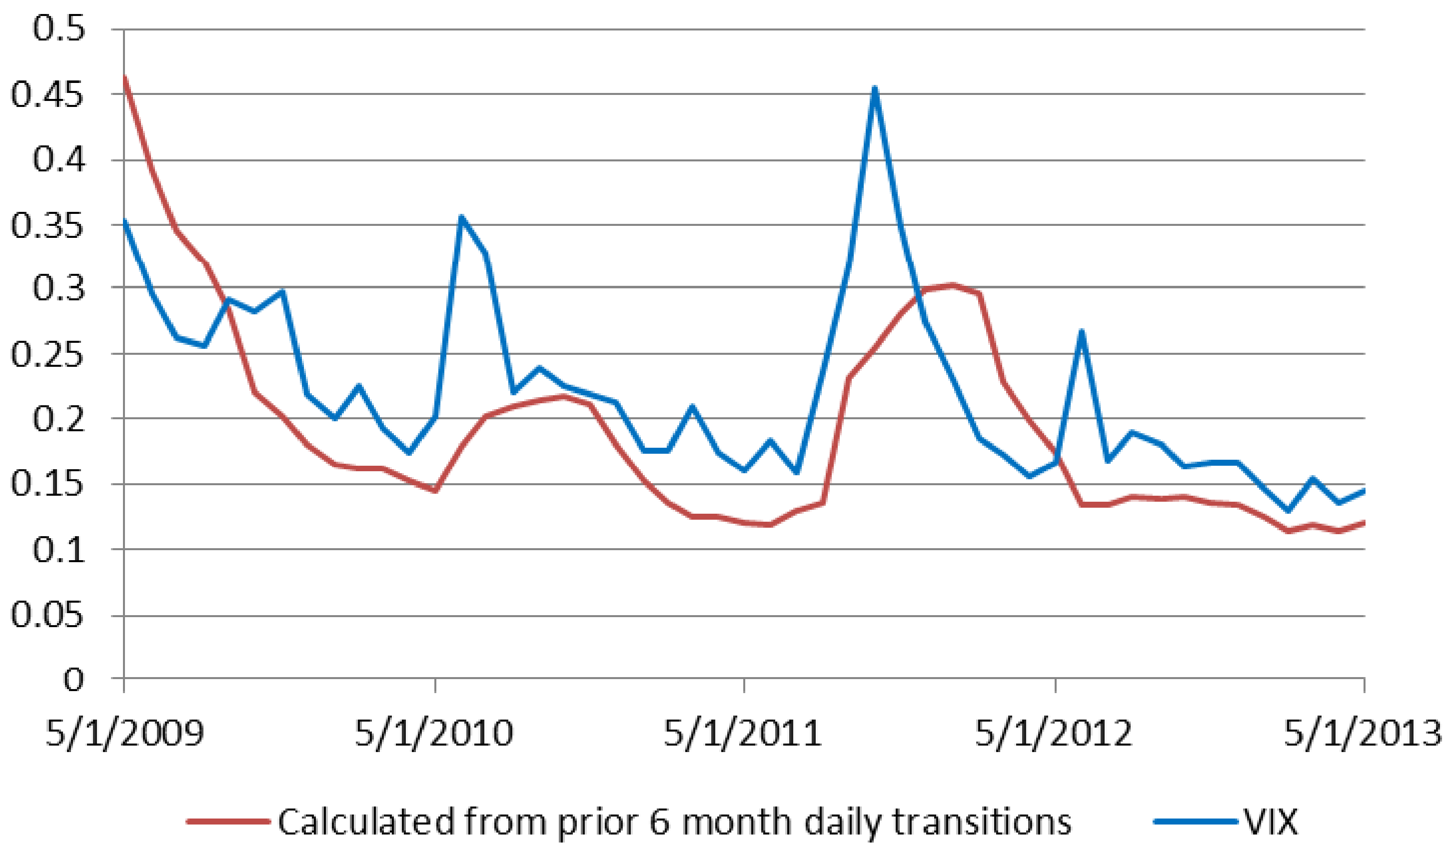

Actual Volatility vs. 3 month lagged VIX
(May 1, 2009 – May 1, 2013)

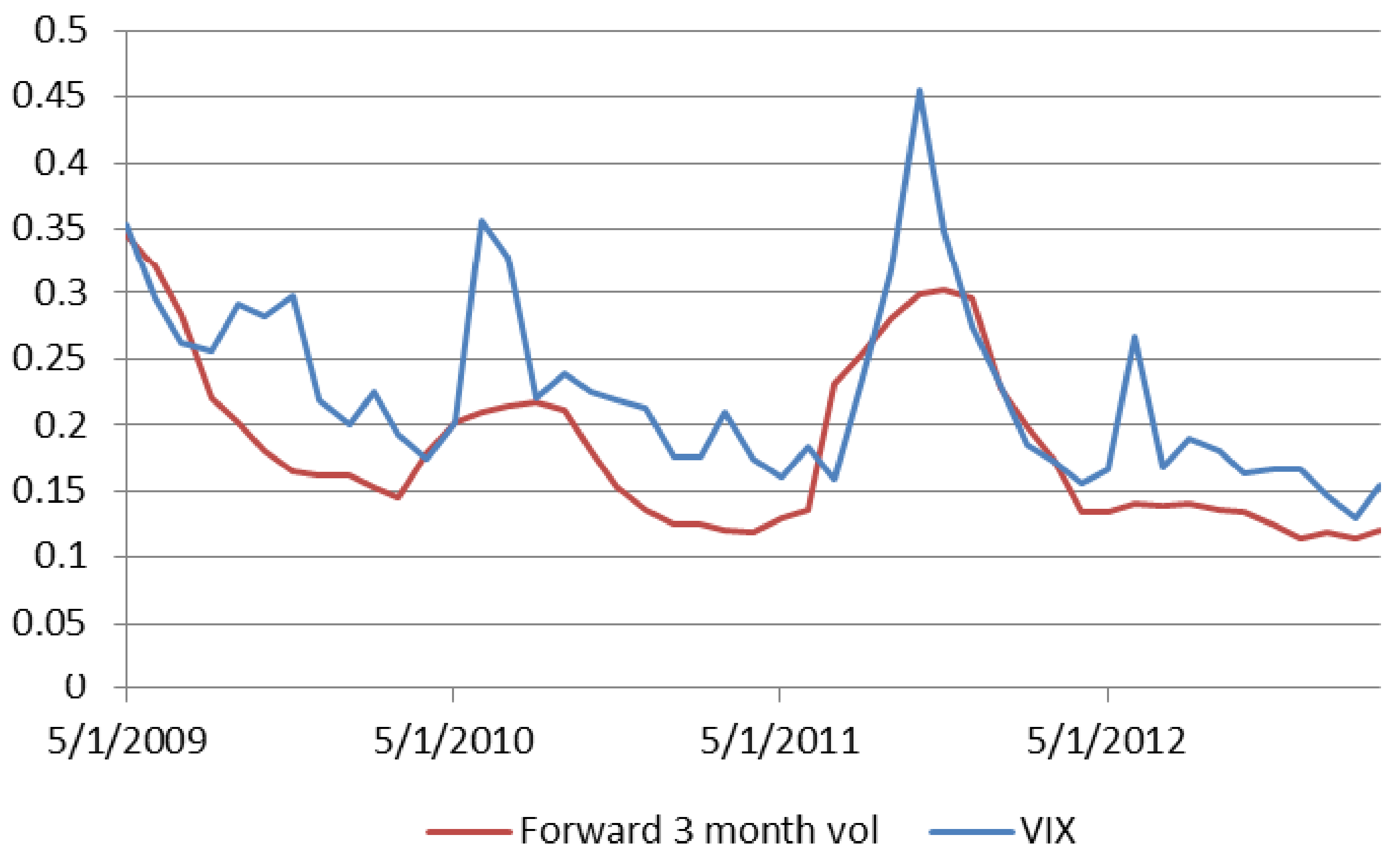

Actual Volatility vs. Recovered Volatility
(May 1, 2009 – May 1, 2013)

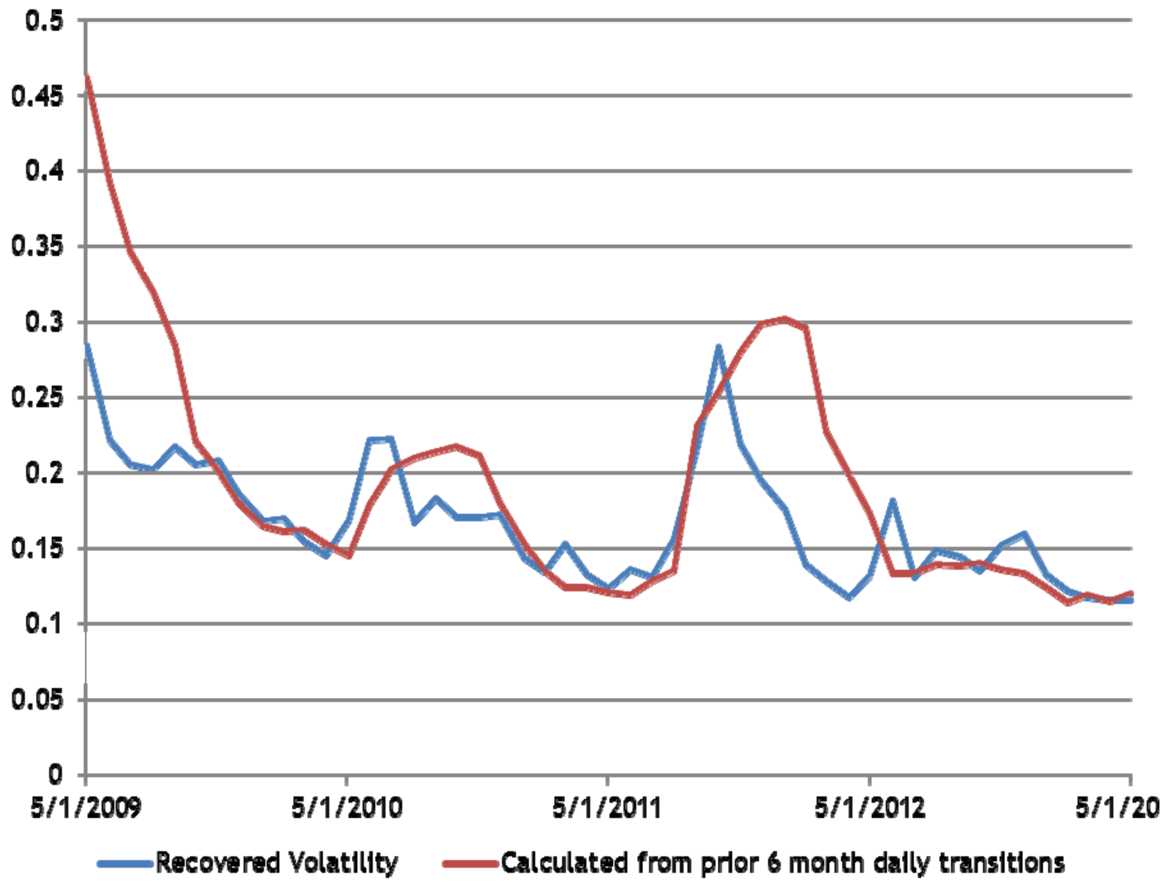

Actual Volatility vs. 3 month lagged Recovered Volatility
(May 1, 2009 – May 1, 2013)

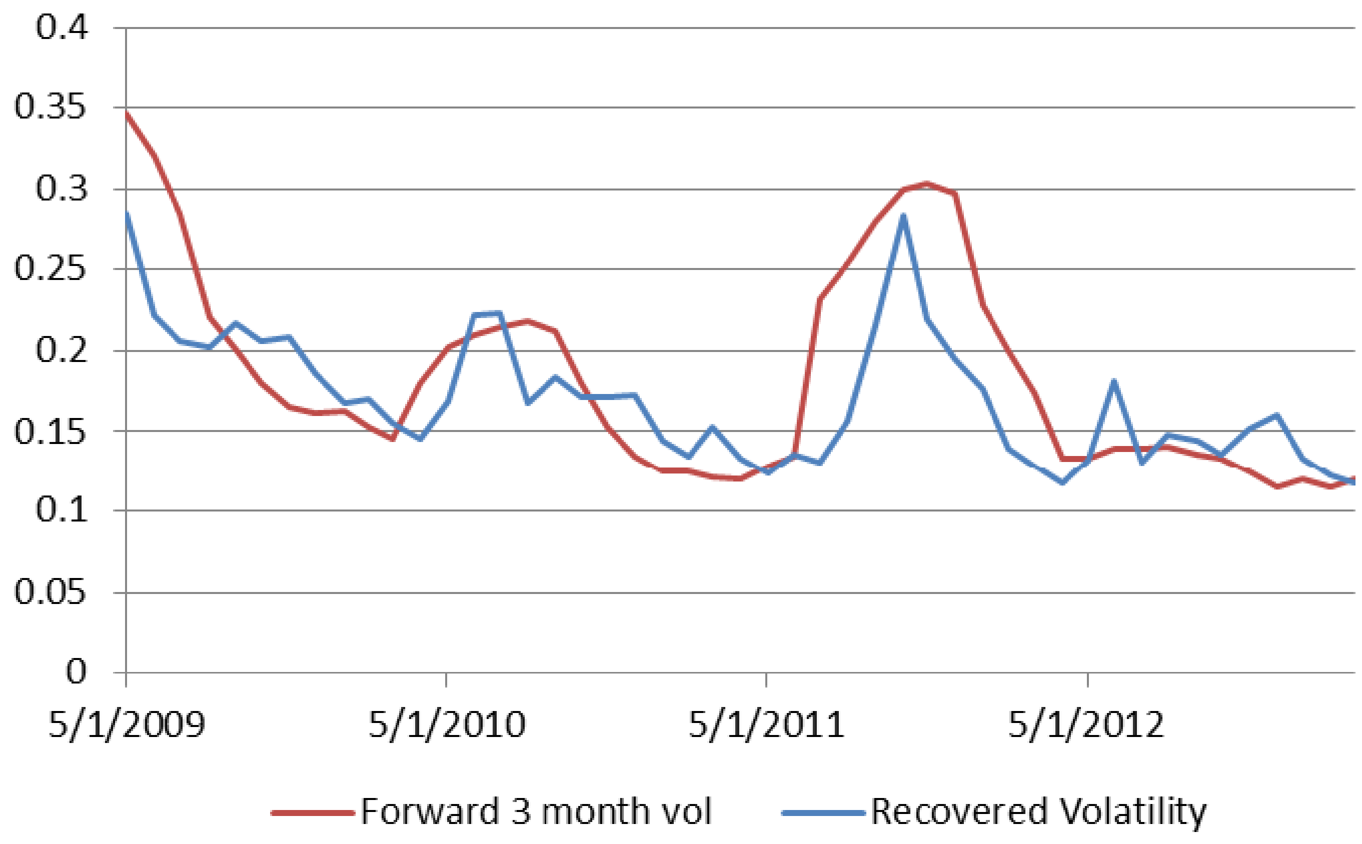

Regression of S&P 500 Monthly Return on Recovered Expected Return
(May 1, 2009 – April 1, 2013)

Regression Statistics	
Multiple R	0.39902
R Square	0.159217
Adjusted R	0.140939
Standard E	0.04286
Observatic	48

ANOVA

	df	SS	MS	F	ignificance F
Regressior	1	0.016002	0.016002	8.710896	0.004964
Residual	46	0.0845	0.001837		
Total	47	0.100502			

	Coefficients	andard Erro	t Stat	P-value	Lower 95%	Upper 95%	ower 95.0%	pper 95.0%
Intercept	-0.03763	0.018369	-2.04857	0.046235	-0.0746	-0.00066	-0.0746	-0.00066
X Variable	10.1396	3.435497	2.951423	0.004964	3.224307	17.0549	3.224307	17.0549

Lagged Recovered Regression Adjusted Expected Return vs. S&P 500 Return

(May 1, 2009 – April 1, 2013)

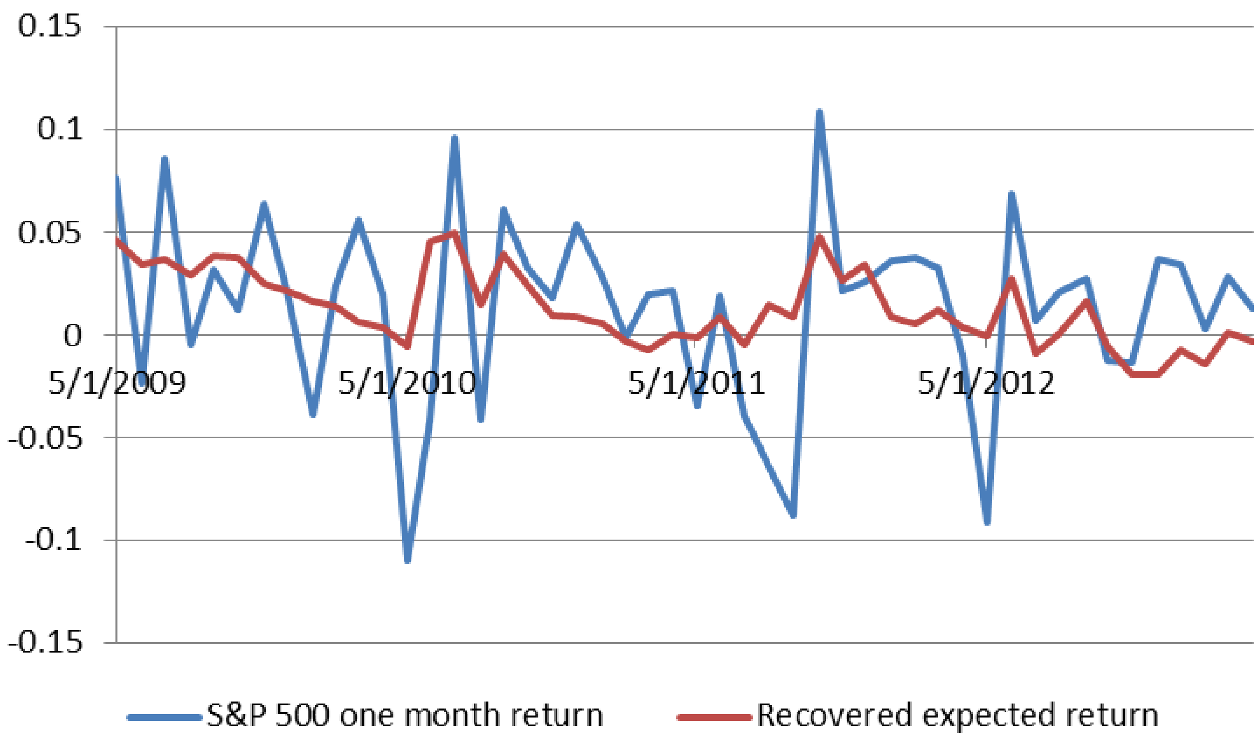

Regression of S&P 500 Quarterly Return on Recovered Quarterly Expected Return

(May 1, 2009 – April 1, 2013)

Regression Statistics	
Multiple R	0.403415
R Square	0.162744
Adjusted R	0.143715
Standard E	0.067543
Observatic	46

ANOVA

	df	*SS*	*MS*	*F*	*ignificance F*
Regressior	1	0.039017	0.039017	8.55262	0.005436
Residual	44	0.20073	0.004562		
Total	45	0.239747			

	Coefficients	*andard Erro*	*t Stat*	*P-value*	*Lower 95%*	*Upper 95%*	*ower 95.0%*	*pper 95.0%*
Intercept	-0.06054	0.035068	-1.72624	0.091322	-0.13121	0.010139	-0.13121	0.010139
X Variable	5.710293	1.95258	2.924486	0.005436	1.775127	9.645459	1.775127	9.645459

Quarterly Lagged Recovered Regression Adjusted Expected Return vs. Quarterly S&P 500 Return

(May 1, 2009 – April 1, 2013)

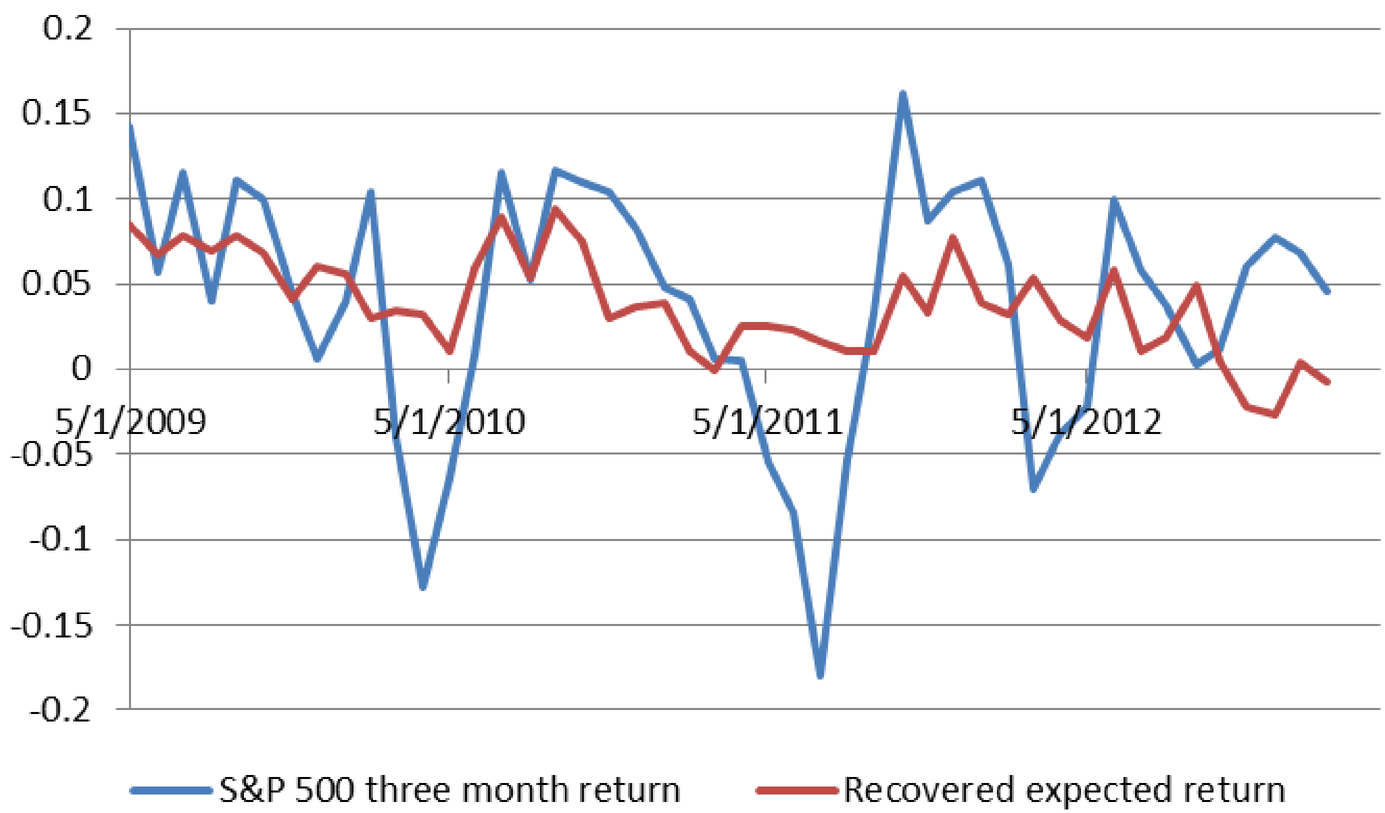

Some Applications and a To Do List

- We should test the method by recovering the distribution at a series of historical dates, and then compare those forecasts with the actual subsequent outcomes
- We should also compare the recovered predictions with the historical distributions and with other economic and capital market factors to find potential hedge and/or leading/lagging indicator relationships
- This should prove valuable for asset allocation and a host of practical problems, e.g., determining the market risk premium
- Extending the analysis to the fixed income markets is already underway by Peter Carr and his coauthors and in joint research by myself and Ian Martin of Stanford
- Publish a monthly report on the current recovered characteristics of the stock return distribution:
  - The forecast equity risk premium
  - The chance of a catastrophe or a boom

Appendix

A Simple Recovery Theorem

- The Recovery Theorem is not much deeper than simply observing that we have the same number of equations as unknowns
- The unknowns are the market transition probabilities, f_{ij} , and the risk aversions, R(j)
- If m is the number of states, then there are a total of $m + m^2$ unknowns (ignoring δ which can also be found)
- The equations are the m^2 pricing equations:

$$p_{ij} = \delta \varphi_{ij} f_{ij} = \delta \left(\frac{R(j)}{R(i)} \right) f_{ij}$$

- And, also, for any state i, the probabilities of going to some state have to add to one:

$$f_{i1} + \cdots + f_{im} = 1$$

- This is a total of $m + m^2$ in the $m + m^2$ unknowns so, as is usual in algebra when the number of equations equals the number of unknowns, we can solve it!

Bounded Eigenfunctions

- In a continuous space the matrix operators are integrals

$$A(g(\cdot)) = e^{-\delta T}\frac{1}{\varphi(x)}\int \varphi(y)g(y)f(x,y)dy$$

Since the kernel is positive, A is a positive linear operator. Notice that

$$A\left(\frac{1}{\varphi}\right) = e^{-\delta T}\frac{1}{\varphi(x)}\int \varphi(y)\frac{1}{\varphi(y)}f(x,y)dy = e^{-\delta T}\frac{1}{\varphi(x)}$$

- **Theorem 1: The operator, A, has a unique positive eigenvalue, $e^{-\delta T}$, associated with all positive eigenvectors bounded both from above and away from zero.**

- **Theorem 2: The kernel is the unique eigenfunction of A that is bounded away from zero and of bounded variation.**

- **Theorem 3 The Continuous State Space Recovery Theorem**

  Under the assumptions outlined above, it is possible to recover both the pricing kernel, φ, and the natural probability density, f(x,y) from the conditional state price density.

A Simple Binomial Example

- If m = 2 we have a simple 2x2 example:

$$\varphi_1 p_{11} = \delta f_1 \varphi_1$$

$$\varphi_1 p_{12} = \delta(1 - f_1)\varphi_2$$

$$\varphi_2 p_{21} = \delta(1 - f_2)\varphi_1$$

and

$$\varphi_2 p_{22} = \delta f_2 \varphi_2$$

- This is four equations in five unknowns, ϕ_1, ϕ_2, f_1, f_2, and δ which, with some manipulation reduces to a quadratic in δ with a unique positive solution

- Without state dependence we would have

$$p_1 = \delta f \varphi_1 \quad and \quad p_2 = \delta(1 - f)\varphi_2$$

two equations in the four unknowns with no unique positive solution