

NBER WORKING PAPER SERIES

THE RECOVERY THEOREM

Stephen A. Ross

Working Paper 17323

<http://www.nber.org/papers/w17323>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

August 2011

The author is a Managing Partner of Ross Farrar, LLC., a Registered Investment Adviser that manages portfolios of derivative securities including options. The data used in the paper were obtained under a non-disclosure agreement with a financial firm. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2011 by Stephen A. Ross. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Recovery Theorem
Stephen A. Ross
NBER Working Paper No. 17323
August 2011
JEL No. E1,G0,G11,G12,G17

ABSTRACT

We can only estimate the distribution of stock returns but we observe the distribution of risk neutral state prices. Risk neutral state prices are the product of risk aversion – the pricing kernel – and the natural probability distribution. The Recovery Theorem enables us to separate these and to determine the market's forecast of returns and the market's risk aversion from state prices alone. Among other things, this allows us to determine the pricing kernel, the market risk premium, the probability of a catastrophe, and to construct model free tests of the efficient market hypothesis.

Stephen A. Ross
MIT Sloan School of Management
100 Main Street, E62-616
Cambridge, MA 02142
and NBER
sross@mit.edu

Introduction

Because financial markets price securities with payoffs extending out in time the hope that they can be used to forecast the future has long fascinated both scholars and practitioners. Nowhere has this been more apparent than with the studies of the term structure of interest rates with its enormous literature devoted to examining the predictive content of forward rates. But with the exception of foreign exchange and some futures markets, a similar line of research has not developed in other markets and, most notably, not in the equity markets.

While we have a rich market in equity options and a well-developed theory of how to use their prices to extract the martingale risk neutral probabilities (see Cox and Ross (1976a, 1976b)), there has been a theoretical hurdle to using them to make forecasts or, for that matter, to speak to issues in the natural world. Risk neutral returns are natural returns that have been ‘risk adjusted’. In the risk neutral measure the expected return on all assets is the risk free rate because the risk neutral measure is the natural measure with the risk premium subtracted out. The risk premium is a function both of risk and of the market’s risk aversion, and to use risk neutral prices to inform about real or natural probabilities we have to know the risk adjustment so we can add it back in. In models with a representative agent this is equivalent to knowing that agent’s utility function and that is not directly observable. Instead, we infer it from fitting or ‘calibrating’ market models.

Furthermore, efforts to empirically measure the aversion to risk have led to more controversy than consensus. For example, measurements of the coefficient of aggregate risk aversion range from 2 or 3 to 500 depending on the model. The data are less helpful than we would like because we have a lengthy history in which U.S. stock returns seemed to have consistently outperformed fixed income returns – the equity premium puzzle (Prescott and Mehra [1985])– and that has given rise to a host of suspect proscriptions for the unwary investor. These conundrums have led some to propose that finance has its equivalent to the dark matter cosmologists posit to explain the behavior of their models for the universe when observables don’t seem to be sufficient. Our dark matter is the very low probability of a catastrophic event and the impact that changes in that perceived probability can have on asset prices (see, e.g., Barro [2006] and Weitzmann [2007]). Apparently, though, such events are not all that remote and five sigma events seem to occur with a frequency that belies their supposed low probability.

When we extract the risk neutral probabilities of such events from the prices of options on the S&P 500, we find the risk neutral probability of, for example, a 25% drop in a month, to be higher than the probability calculated from observed stock returns. But, since the risk neutral probabilities are the natural probabilities adjusted for the risk premium, either the market forecasts a higher probability of a stock decline than occurred historically or the market requires a very high risk premium to insure against a decline. Without knowing which, it is impossible to separate the two out and find the market’s forecast of the event probability.

Finding the market's forecast for returns is important for other reasons as well. The natural expected return of a strategy depends on the risk premium for that strategy and, consequently, it has long been argued that any tests of efficient market hypotheses are simultaneously, tests of a particular asset pricing model and of the efficient market hypothesis (Fama [1970]). But if we knew the kernel we could estimate how variable the risk premium is (see Ross [2005]), and a bound on the variability of the kernel would limit how predictable a model for returns could be and still not violate efficient markets. In other words, it would provide a model free test of the efficient markets hypothesis.

A related issue is the inability to find the current market forecast of the expected return on equities. Unable to read this off of prices as we do with forward rates, we are left to using historical returns and resorting to opinion polls of economists and investors - asking them to reveal their estimated risk premiums. It certainly doesn't seem that we can derive the risk premium directly from option prices because by pricing one asset - the derivative - in terms of another, the underlying, the elusive risk premium doesn't appear in the resulting formula.

But, in fact, all is not quite so hopeless. While quite different, our results are in the spirit of Dybvig and Rogers [1997], who showed that if stock returns follow a recombining tree (or diffusion) then from observing an agent's portfolio choice along a single path we can reconstruct the agent's utility function. We borrow their nomenclature and call our results recovery theorems as well. Section 1 provides an overview of previous approaches to these problems. It estimates the pricing kernel by dividing the risk neutral martingale density derived from the market by the estimated natural distribution. As an alternative, it calibrates some standard models for preferences and uses the resulting marginal rates of substitution for the kernel to derive the natural distribution. Section 2 develops the basic analytic framework tying the natural probabilities to the risk neutral probabilities and proves the Recovery Theorem. Here we introduce the key idea of the paper, a technique to estimate the natural probability of asset returns and the market's risk aversion, the kernel, from the state price transition process alone. Section 3 derives a second recovery theorem, the Multinomial Recovery Theorem, which offers an alternative route for recovering the natural distribution for binomial and multinomial processes. Section 4 examines the application of these results to some examples and examines some extensions. Section 5 estimates the state price densities at different horizons from the S&P 500 option prices on a randomly chosen recent date, April 27, 2011. As a prelude we tidy up some loose ends in the literature on no arbitrage pricing of options. Section 6 estimates the state price transition matrix and applies the Recovery Theorem to derive the kernel and the natural probability distribution. We compare the model's estimate of the natural probability with the histogram of historical stock returns. In particular, we shed some light on the dark matter of economics by highlighting the difference between the odds of a catastrophe as derived from observed state prices with that obtained from historical data. The analysis of Sections 5 and 6 is meant to be illustrative and is far from the much needed empirical analysis, but it provides the first estimate of the natural density of stock returns. Section 7 derives a model free test of efficient market hypotheses. Section 8 concludes and summarizes the paper and points to some future research directions.

Section 1: Estimating Natural Probabilities, Risk Aversion and Risk Neutral Probabilities

Consider a one period world with asset payoffs $x(\theta)$ at time T , contingent on the realization of a state of nature, $\theta \in \Omega$. From the Fundamental Theorem of Asset Pricing (see Dybvig and Ross [1987, 2003]), no arbitrage ('NA') implies the existence of positive state space prices, i.e., Arrow Debreu contingent claims prices, $p(\theta)$ (or allowing for lumpy states, the distribution function, $P(\theta)$), paying \$1 in state θ and nothing in any other states. If the market is complete, then these state prices are unique. The current value of the asset is given by

$$p = \int x(\theta) dP(\theta)$$

Since the sum of the contingent claims prices is the current value of a dollar for sure in the future, letting r denote the riskless rate we can rewrite this in the familiar forms

$$\begin{aligned} p &= \int x(\theta) dP(\theta) \\ &= \left(\int dP(\theta) \right) \int x(\theta) \frac{dP(\theta)}{\int dP(\theta)} \equiv e^{-rT} \int x(\theta) \pi^*(\theta) \equiv e^{-rT} E^*[x(\theta)] \\ &= e^{-rT} E[x(\theta)\varphi(\theta)] \end{aligned}$$

where an asterisk denotes the expectation in the martingale measure and where the pricing kernel, i.e., the state price/probability $\mathbb{Q}(\theta)$ is the Radon-Nikodym derivative of $P(\theta)$ with respect to the natural measure which we will denote as $F(\theta)$. With continuous distributions, $\mathbb{Q}(\theta) = p(\theta)/f(\theta)$ where $f(\theta)$ is the natural probability. The risk neutral probabilities, are given by $\pi^*(\theta) = p(\theta)/\int p(\theta) = e^{rT} p(\theta)$.

Equivalently we can rewrite this in terms of (gross) returns,

$$E^*[R_x] = E[R_x \varphi(\theta)] = E^* \left[\frac{x(\theta)}{p} \right] = e^{rT}$$

Notice, that the interest rate, r , could be dependent on the current state, a possibility we will allow below.

Denoting by Y , the stock market index value at time T , and the current value of the stock market index as S , in a one period model the absence of arbitrage implies the existence of a monotone increasing concave utility function for which the optimum portfolio problem has a solution (see Dybvig and Ross [1987, 2003]). In addition, if individuals maximize expected utility, then in a complete market, there will be a unique utility function for the representative agent, $U(Y)$, that is maximized at Y such that

$$\varphi(S, Y) = \delta \frac{U'(Y)}{U'(S)}$$

where δ is the subjective time discount factor between today and time T. Indeed, the assumed existence of such a representative agent is the foundation for the intuition of a market risk aversion.

This equation tells us that the state price density, φ , relating risk neutral probabilities and natural probabilities is the marginal rate of substitution for the payoffs in a one period world. This is the barrier between relating the one to the other and why we cannot simply use risk neutral probabilities as though they were natural probabilities.

Oddly, it is often said that while we observe the natural probabilities from the data on stock returns we don't observe the martingale probabilities. In fact, the opposite is true in the financial markets; the martingale risk neutral probability density $\pi^*(\theta)$ is observable and it is the natural probability measure, $f(\theta)$ that must be estimated. Figure 1 illustrates the estimated $\pi^*(\theta)$ derived from six month option prices on March 15, 2011 (in Section 5 we show how this was done). Figure 2 displays an estimated lognormal density for actual stock returns and Figure 3 displays the resulting pricing kernel, $p(S)/f(S)$ (S at time 0 is set at 1).

The ridiculously high value for the kernel at low stock prices has two possible interpretations. First, it could mean that the estimated distribution of stock returns has too small a left tail and that the actual probabilities used by the market are higher. Alternatively, it could mean that these probabilities are correct but that the market assigns a very high degree of risk aversion to these outcomes.

To improve on this let's estimate the natural distribution not by fitting it to a given parametric density function such as the log normal but, rather, by letting the 'data speak for itself and using a bootstrap to estimate the return density. Figure 4 shows the result from bootstrapping (with replacement) 5000 six month return samples from monthly returns in the fifty year period, January, 1960 to December, 2010. Figure 5 displays the resulting pricing kernel.

While we can at least display the density, this is only because we've learned our lesson and plotted the kernel in the range where the histogram of actual returns is positive. Of course this is cheating because the martingale measure gives positive probability to the tails where there are no histogram observations and we know that the martingale measure is absolutely continuous with respect to the natural measure. Again, the data is telling us that there are fat tails with high probabilities of bad outcomes or with high risk aversion or both. Since all of the action, i.e., the real volatility and pricing impact of the kernel is in the left tail this approach clearly misses the point. The resulting computed volatility of the kernel, .0058, is unreasonably low as well precisely because of this truncation.

Let's try a different approach. Rather than estimating the natural distribution of stock returns, let's posit a utility function, derive the kernel as the marginal rates of substitution, and divide the martingale measure by the marginal rates of substitution to obtain an estimate of the natural distribution. The most common functional form for the utility function is a constant relative risk aversion,

$$U(Y) = \frac{Y^{1-R}}{1-R}$$

with the associated kernel,

$$\varphi(S, Y) = \frac{U'(Y)}{U'(S)} = \frac{Y^{-R}}{S^{-R}} = \left(\frac{Y}{S}\right)^{-R}$$

For simplicity, we have omitted the subjective rate of time discount. Figures 6 and 7 plot the kernel for three choices of risk aversion, $R = 0$, risk neutrality, $R = 0.5$ and $R = 5.0$. Figure 8 plots the respective distributions of the natural probabilities. Notice that with risk aversion, $R = 0$, the natural density is just the risk neutral density.

Not surprisingly, the lower the risk aversion the tighter the bound on the kernel and we have that $E[\mathbb{Q}^2] = 1, 1.03$ and 56 for $R = 0, 0.5,$ and 5.0 respectively. Without comparing this with the measured empirical distribution, though, there is no obvious way to choose amongst these estimates. Furthermore, we know that if we simply follow the usual path of estimating or calibrating the kernel marginal rates of substitution to return data we are just going down the path where whatever estimate we find will have some insensibility, e.g., too high a risk premium, too low a risk aversion coefficient or too much required volatility

What have we learned from these exercises? Contrary to what we might have thought and what we will illustrate below, it is the martingale measure that is directly observable and the natural measure that can only be estimated with unacceptable errors for our purposes. In the next sections we will take a closer look at the theory and show the full range of information that we can glean from the risk neutral measure.

Section 2: The Recovery Theorem

In this section we set up the basic analytic model and derive the Recovery Theorem, which enables us to find the natural measure and the pricing kernel from state prices alone.

We will continue to assume that we have observed the martingale measure (for example, by observing a full spectrum of option prices). That means that we have not only observed the current state prices but, also, the martingale transition probabilities. Let x denote the current state and y a state one period forward. We assume that this is a full description of the state of nature including the stock price itself and other information that is pertinent to the future evolution of the stock market index, thus the stock price can be written as $S(x)$. From the forward equation for the martingale probabilities

$$Q(x, y, T) = \int Q(x, z, t)Q(z, y, T - t)dz$$

where $Q(x, y, \tau)$ is the forward probability transition function for going from state x to state y in τ periods and where the integration is over the intermediate state at time t, z . Notice that we have made the transition a function independent of calendar time and a function only of the time interval.

This is a very general framework and allows for many interpretations. For example, the state could be composed of parameters that describe the motion of the process, e.g., the volatility of returns, v , as well as the current stock price, S , i.e., $x = (S, v)$. If the distribution of martingale returns is determined only by the volatility, then a transition could be written as a move from $x = (S, v)$ to $y = (S(1+R), z)$ where R is the rate of return and

$$Q(x, y, t) = Q((S, v), (S(1 + R), z), t)$$

To simplify notation we will use state prices rather than the martingale probabilities so that we won't have to be continually correcting for the interest factor. Defining the state price matrix as

$$p(x, y, t, T) \equiv e^{-r_x T} Q(x, y, t, T)$$

and, assuming a time homogeneous process, period by period we have,

$$p(x, y) = e^{-r_x} Q(x, y)$$

Notice that the risk free rate is allowed to be state dependent. This is appropriate and if we were to fix it across states we will see below that the model would degenerate.

Letting f denote the natural (time homogeneous) transition density, the kernel in this framework is defined as the price per unit of probability,

$$\varphi(x, y) = \frac{p(x, y)}{f(x, y)}$$

Specializing this to an intertemporal model with additively time separable preferences and a constant discount factor, δ , and letting c denote consumption at time t as a function of the state, the kernel can be written as

$$\varphi(x, y) = \frac{p(x, y)}{f(x, y)} = \frac{\delta U'(c(y))}{U'(c(x))}$$

Since this equation is simply the first order condition for the optimum for a representative agent,

$$\begin{aligned} \max_{\{c(x), c(y)\}} \{ & U(c(x)) + \delta \int U(c(y)) f(x, y) dz, \\ \text{s. t. } & c(x) + \int c(y) p(x, y) dz = w \end{aligned}$$

the equation for the kernel is the equilibrium solution for an economy with complete markets in which, for example, consumption is exogenous and prices are defined by the first order condition for the optimum. In a multiperiod model with complete markets and state independent intertemporally additive separable utility, there is a unique representative agent utility function that satisfies the above optimum condition and determines the kernel as a function of aggregate consumption (see Dybvig and Ross [1987, 2003]). Notice, too, that we don't have state dependent utility; the pricing kernel depends only on the marginal rate of substitution between future and the current consumption. The existence of such a representative agent will be a maintained assumption of our analysis below.

Rewriting the equilibrium equation with additively time separable utility we have,

$$U'(c(x))p(x, y) = \delta U'(c(y))f(x, y)$$

To gain some insight into this equation and to position the apparatus for empirical work it's useful to look at a discrete state model

$$U'_i p_{ij} = \delta U'_j f_{ij}$$

where

$$U'_i \equiv U'(c(i))$$

and writing this in terms of the kernel

$$\varphi_i \equiv \varphi(1, i) = \delta(U'_i/U'_1)$$

We can interpret the marginal utilities, U_i' , as the marginal utility of consumption in state i , $c(i)$, and while it is monotone declining in consumption it need not be monotone declining in the asset value, $S(i)$. In practice, we will define the states from the motion of the stock value, so that the kernel is the projection of the kernel across the broader state space onto the more limited space defined by the filtration of the asset price.

Rewriting the state equations in matrix form we have

$$DP = \delta FD$$

where P is an $m \times m$ matrix, F is the $m \times m$ matrix of the natural probabilities and D is the diagonal matrix with the marginal utilities on the diagonal

$$D = \begin{bmatrix} U'_1 & 0 & 0 \\ 0 & U'_i & 0 \\ 0 & 0 & U'_m \end{bmatrix} = \begin{bmatrix} \varphi_1 & 0 & 0 \\ 0 & \varphi_i & 0 \\ 0 & 0 & \varphi_m \end{bmatrix}$$

With a discrete or compact state space for prices we will have to make sure that the model doesn't admit of arbitrage. In the model with exogenous consumption whose motion is governed by the specified state transition density the absence of arbitrage is a simple consequence of an equilibrium with positive state prices which assures that the carrying cost net of the dividend compensates for any position that attempts to profit from the rise out of the lowest asset value or the decline from the highest value.

Returning to our analysis, keep in mind that we observe the martingale prices, P , and our objective is to see what, if anything, we can infer about the natural measure, F , and the pricing kernel, i.e., the marginal rates of substitution. Solving for F as a function of P ,

$$F = \left(\frac{1}{\delta}\right)DPD^{-1}$$

Clearly if we knew D , we would know F . It appears that we only have m^2 equations in the m^2 unknown probabilities, the m marginal utilities, and the discount rate, δ , and this appears to be the current state of thought on this matter. We know the risk neutral measure but without the marginal rates of substitution across the states, i.e., the risk adjustment, there appears to be no way to close the system and solve for the natural measure, F . Fortunately, though, this must satisfy an additional set of m constraints, namely that F is a stochastic matrix whose row sums to 1.

$$Fe = e$$

where e is the vector with 1 in all the entries.

Using this condition we have

$$Fe = \left(\frac{1}{\delta}\right)DPD^{-1}e = e$$

or

$$Pz = \delta z$$

where

$$z = D^{-1}e$$

This is a characteristic root problem and offers some hope that the solution set will be discrete and not an arbitrary large linear space. The theorem below verifies this intuition and provides us with a powerful result.

Theorem 1 – The Recovery Theorem

In a world with a representative agent, if the pricing matrix, P , is positive or irreducible, then there exists a unique (positive) solution to the problem of finding F , the discount rate, and the pricing kernel. That is for any given set of state prices there is one and only one corresponding natural measure and, therefore, a unique pricing kernel, \mathbb{Q} . If P has a single absorbing state but the matrix of the remaining $m-1$ states is irreducible, then it has a unique positive solution for states other than for the absorbing state and if the diagonal entry of the absorbing state is greater than the positive characteristic root of the other states, then there is a unique strictly positive solution.

Proof:

Existence follows immediately from the fact that P is assumed to be generated from F and D as shown above, but as a byproduct of our solution we will show existence even if this were not the case. The problem of solving for F is equivalent to finding the characteristic roots (eigenvalues) and characteristic vectors (eigenvectors) of the matrix of state prices, P . If we know δ and z such that

$$Pz = \delta z$$

then the kernel can be found from $z = D^{-1}e$.

Since P is the matrix of state prices it is a nonnegative matrix. Suppose first that it is strictly positive. Since the martingale measure is absolutely continuous with respect to the natural measure, there are no states that are unattainable and could have zero prices. The Perron-Frobenius Theorem (see Meyer [2000]) tells us that all such matrices have a unique positive characteristic vector, z , and an associated positive characteristic root, λ , and that there are no other strictly positive characteristic vectors. Since P satisfies

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

we know that a positive D must exist that yields P from some given F hence, we have now found it. The characteristic root $\lambda = \delta$ is the subjective rate of time discount.

Now suppose that P is nonnegative and irreducible but not positive, i.e., suppose that there are some zero entries. A zero in the ij entry means that state j cannot be attained from state i in a single step, but the matrix is irreducible if all states are attainable from each other states in n steps. The current price for this state, occurring n periods out, would thus be positive. In this case the Perron-Frobenius Theorem holds and as before there is a unique positive characteristic vector associated with a unique positive characteristic root.

Lastly, suppose that one state has zeros at all points except on the diagonal. This could occur very naturally if the most calamitous state, say state 1, was absorbing and once entered it cannot be left. In this case the entry $p_{11} > 0$ and the remaining entries in the first row are zero. As is well known from the theory of Markov chains, P is a reducible matrix since for all n , the entries in the first row of P^n other than the first will remain zero. If P is reducible then there is not a lot known in general about the characteristic roots, however, if we restrict ourselves to the case described above where the worst case is the single absorbing state and the rest of the entries of P are positive we do have a result.

Partition P as follows where B is $(m-1) \times (m-1)$, and C is $(m-1) \times 1$,

$$P = \begin{bmatrix} p_{11} & 0 \\ C & B \end{bmatrix}$$

and partition the possible characteristic vector as

$$\begin{bmatrix} y \\ x \end{bmatrix}$$

where y is a scalar and x is dimension $m-1$.

To be a characteristic vector we require that

$$P \begin{bmatrix} y \\ x \end{bmatrix} = \begin{bmatrix} p_{11}y \\ yC + Bx \end{bmatrix} = \rho \begin{bmatrix} y \\ x \end{bmatrix}$$

If we set $y = 0$, then a characteristic vector is given by the solution to $Bx = \lambda x$ which is positive and unique since B is irreducible. To find a strictly positive solution, without loss of generality scale $y = 1$. Now we must solve the system

$$C + Bx = p_{11}x$$

If $p_{11} = \lambda$ this generally has no solution; the characteristic matrix

$$[I - \lambda B]$$

is singular and the system can be solved only if C is in its span.

In the general case where $p_{11} \neq \lambda$, the unique solution is given by

$$x = -[B - p_{11}I]^{-1}C$$

Since λ is the maximal characteristic root, the inverse is negative for $p_{11} > \lambda$ and since $C > 0$ this assures that x is strictly positive.

□

Letting x denote the unique positive characteristic vector with root λ , we can solve for the kernel as

$$U'(c(i)) = \left(\frac{1}{\delta}\right) \varphi_i = d_{ii} = \frac{1}{x_i}$$

Since x can be arbitrarily scaled without changing the marginal rates of substitution we can choose a reference state, c for the current state, and divide each entry by this marginal utility. From our previous analysis,

$$F = \left(\frac{1}{\delta}\right) DPD^{-1}$$

and

$$f_{ij} = \left(\frac{1}{\delta}\right) \frac{\varphi_i}{\varphi_j} = \left(\frac{1}{\delta}\right) \frac{U'_i}{U'_j} p_{ij} = \left(\frac{1}{\lambda}\right) \frac{x_j}{x_i} p_{ij}$$

Corollary 1

The subjective rate of discount, δ , is bounded above by the largest interest factor.

Proof:

From The Recovery Theorem the subjective rate of discount, δ , is the positive character root of the price transition matrix, P . From the Perron-Frobenius Theorem (see Meyer [2000]) this root is bounded above by the maximum row sum of P . Since the elements of P are the pure contingent claim state prices, it follows that the row sums are the interest factors and that the maximum row sum is the maximum interest factor.

□

Now let's turn to the case where the riskless rate is the same in all states.

Theorem 2

If the riskless rate is state independent then the unique natural density associated with a given set of risk neutral prices is the martingale density itself, i.e., pricing is risk neutral.

Proof:

In this case we have

$$Pe = \gamma e$$

where γ is the interest factor. It follows that $Q = (1/\gamma)P$ is the risk neutral martingale probability matrix and, as such, e is its unique positive characteristic vector and 1 is its characteristic root.

From Theorem 1

$$F = \left(\frac{1}{\gamma}\right)P = Q$$

□

Given the apparent ease of creating intertemporal models satisfying the usual assumptions without risk neutrality this result seems a bit odd, but it's a consequence of having a finite irreducible process for state transition. Apparently when we extend the recovery result to multinomial processes that are unbounded this is no longer the case.

Before going on to implement these results, there is a simple extension of this approach that appears not to be well known and is of interest in its own right.

Theorem 3

The risk neutral martingale density for consumption and the natural density for consumption have the single crossing property and the natural density stochastically dominates the risk neutral density. Equivalently, in a one period world, the market natural density stochastically dominates the risk neutral density.

Proof:

From

$$\varphi(x, y) = \frac{p(x, y)}{f(x, y)} = \frac{\delta U'(c(y))}{U'(c(x))}$$

we know that φ is declining in $c(y)$. Since both densities integrate to one and since φ exceeds $1/\delta$ for $c(y) < c(x)$, it follows that $p > f$ for $c(y) < c(y^*)$ where $\delta U'(c^*(y^*)) = U'(c(x))$ and $p < f$ for $c > c^*$. This is the single crossing property and verifies that f stochastically dominates p . In a single period model, terminal wealth and consumption are the same.

□

Corollary 2

In a one period world the market displays a risk premium, i.e., the expected return on the asset is greater than the riskless rate.

Proof:

In a one period world consumption coincides with the value of the market. From stochastic dominance at any future date, T ,

$$S^* \sim S - Z + \epsilon$$

where the asterisk denotes the price in the martingale risk neutral measure, Z is strictly nonnegative and ϵ is a mean zero error term. Taking expectations we have

$$E[S] = r + E[z] > r.$$

□

Before leaving this section we should note that the continuous state space problem has the same form as that for the Recovery Theorem,

$$p(x, y) = \delta \frac{U'(c(y))}{U'(c(x))} f(x, y)$$

and solving for f

$$f(x, y) = \left(\frac{1}{\delta}\right) \frac{U'(c(x))}{U'(c(y))} p(x, y)$$

The resulting integral equation analogous to the matrix equation above that f must satisfy is given by

$$\int f(x, y) dy = 1$$

implying that

$$\int \frac{1}{U'(c(y))} p(x, y) dy = \delta \frac{1}{U'(c(x))}$$

This is an integral form of the matrix equation to which we applied the Perron-Frobenius Theorem to demonstrate that it had a unique positive solution. Extensions of the Perron Frobenius Theorem to positive integral operators are available which would allow us to find a unique (up to a scalar multiplier) positive solution for $U'(S)$ and a unique positive root δ . There is, however, more work to be done to satisfy the requirements of these extensions, e.g., behavior in the tails, and undertaking that would take us beyond the scope of this paper. Here we have developed the central argument and its intuitions and derived a sufficiently powerful framework for the subsequent empirical analysis. The next section examines the application of this framework to some examples created to illustrate how it can be applied.

Section 4: A Binomial and Multinomial Recovery Theorem

j

While we will see below that the Recovery Theorem can be applied to a binomial or multinomial process, these processes are so ubiquitous in finance (see Cox, Ross Rubinstein [1979]), that it is useful to look at them separately. Throughout this analysis the underlying metaphorical model is a tree of height H that grows exogenously and bears exogenous fruit, ‘dividends’, that is wholly consumed. Tree growth is governed by a multinomial process and the state of the economy is $\langle H, i \rangle$, $i = 1, \dots, m$. The multinomial process is state dependent and the tree grows to $a_j H$ with probability f_{ij} . In every period the tree pays a consumption dividend kH where k is a constant. Notice that the state only determines the growth rate and the current dividend depends only on the height of the tree and not of the state. The value of the tree, the market value of the economy’s assets, is given by $S = S(H, i)$. Since tree height and, therefore, consumption follow a multinomial process, S also follows a multinomial, but, in general, jump sizes will change with the state.

The initial marginal utility of consumption is $U'(kH)$ which is independent of the state variable, i , and, without loss of generality we can set $U'(kH) = 1$. The equilibrium equations are

$$p_{ij}(H) = \delta U'(ka_j H) f_{ij}$$

or, in terms of the undiscounted kernel,

$$p_{ij}(H) = \delta \varphi_j f_{ij}$$

In matrix notation,

$$P = \delta F D$$

$$F = \left(\frac{1}{\delta}\right) P D^{-1}$$

and since F is a stochastic matrix

$$F e = \left(\frac{1}{\delta}\right) P D^{-1} e = e$$

or

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

Assuming P is of full rank, this solves for the kernel, D , as

$$\left(\frac{1}{\delta}\right) D^{-1} e = P^{-1} e$$

Now F is recovered as

$$F = \left(\frac{1}{\delta}\right) PD^{-1}$$

Proceeding along the tree in the same fashion, node by node, we can recover the distribution at all future nodes.

Notice that this analysis only determines δD and not δ and \mathbb{Q} separately. By taking advantage of the recombining feature of the process, though, we can recover δ and \mathbb{Q} separately. For simplicity, consider a binomial process with jumps to a or b . The binomial is recurrent, i.e., it eventually returns arbitrarily close to any starting position. Along with the transition matrix being of full rank, this was one of the key assumptions of the Recovery Theorem. For a binomial, with jumps of a and b , if we were to expand it into an infinite matrix the matrix would have only two nonzero elements in any row, and at a particular node we would only see the marginal price densities at that node. To observe the transition matrix we want to return to that node from a different path. For example, if the current stock price is 1 and there is no exact path that returns to 1, then we can get arbitrarily close to 1 along a path where the number of a steps, i and the number of down steps, $n - i$, satisfy

$$\frac{i}{n - i} \rightarrow -\frac{\log b}{\log a}$$

for large n .

Sparing the obvious continuity analysis, we will simply assume that the binomial recurs in two steps, i.e., $ab = 1$. That implies that it must satisfy the further state equation when it returns to H from having gone to aH ,

$$p_{ab}(aH) = \delta \left(\frac{U'(kH)}{U'(kaH)} \right) f_{ab} = \delta \left(\frac{1}{\varphi_a} \right) f_{ab}$$

Since we can recover $\delta \mathbb{Q}_a = (P^{-1})_a e$, we can solve separately for δ and \mathbb{Q}_a and, more generally, for δ and \mathbb{Q} . The analysis is similar for the general multinomial case.

If the riskless rate is state independent, then P has identical row sums and if it is of full rank, then, as with the first Recovery Theorem, we must have risk neutrality. To see this, let

$$Pe = \gamma e$$

Hence

$$\left(\frac{1}{\delta}\right) D^{-1}e = P^{-1}e = \left(\frac{1}{\gamma}\right) e ,$$

all the marginal utilities are identical and the natural probabilities equal the martingale probabilities.

If P is not of full rank, by the maintained assumption we know that there is a solution to

$$Fe = PD^{-1}e = e$$

In general, though, there is a linear subspace of potential solutions with dimension equal to the rank of P and, not surprisingly, while we can restrict the range of potential solutions, we cannot uniquely recover the kernel and the probability matrix. In particular, if x is a solution to $Px = e$, and y is orthogonal to P , then, $x + y$ is also a solution. In the special case where the interest rate is state independent, though, even if the matrix is of less than full rank than it's easy to see from the above that risk neutrality is one of the potential solutions. We summarize these results in the following theorem.

Theorem 4 - The Multinomial Recovery Theorem

Under the assumed conditions on the process, the kernel, the transition probability matrix and the subjective rate of discount of a binomial (multinomial) process can be recovered at each node from a full rank state price transition matrix alone. If the state prices are independent of the state, then the kernel must be risk neutral. If the transition matrix is of less than full rank, then we can restrict the potential solutions, but we cannot recover uniquely.

Proof:

See above.

□

These results are, of course, implicit in the much studied work on the binomial model, but perhaps because they haven't been the focus of study or perhaps because of the assumption that because of the arbitrary choice of a kernel it would not be possible to separate risk aversion from probabilities, they appear to have gone unnoticed. Which Recovery Theorem should be used in any particular case depends on the exact circumstances.

Constant Relative Risk Aversion

An alternative approach to recovery that allows for more flexibility on the process is to make an assumption on the functional form of the kernel. Suppose, for example, that we make the common assumption of a constant relative risk aversion utility function and that we specialize the model to a binomial with tree growth of a or b , $a > b$. State prices are given by

$$p_{xy}(H) = \varphi(kH, kyH)\delta f_{xy}$$

Hence,

$$S(a, H) = p_{aa}(H)[S(a, aH) + kaH] + p_{ab}(H)[S(b, bH) + kbH]$$

and

$$S(b, H) = p_{ba}(H)[S(a, aH) + kaH] + p_{bb}(H)[S(b, bH) + kbH]$$

Assuming constant relative risk aversion this system becomes linear,

$$\varphi(x, y) = \left(\frac{y}{x}\right)^{-R}$$

with the linear solution

$$S = \gamma_x H$$

where

$$\begin{pmatrix} \gamma_a \\ \gamma_b \end{pmatrix} = \begin{bmatrix} 1 - \delta f_a a^{1-R} & -\delta(1 - f_a)b^{1-R} \\ -\delta(1 - f_b)a^{1-R} & 1 - \delta f_b b^{1-R} \end{bmatrix}^{-1} \begin{pmatrix} \delta f_a k a^{1-R} + \delta(1 - f_a)k b^{1-R} \\ \delta f_b k b^{1-R} + \delta(1 - f_b)k a^{1-R} \end{pmatrix}$$

Thus the stock value S follows a binomial process and at the next step takes on the values $S(a, aH)$ or $S(b, bH)$ depending on the current state and the transition,

$$S(a, H) = \gamma_a H \rightarrow \gamma_a aH = S(a, aH) \text{ or } \gamma_b bH = S(b, bH)$$

$$S(b, H) = \gamma_b H \rightarrow \gamma_a aH = S(a, aH) \text{ or } \gamma_b bH = S(b, bH)$$

Notice, though, that even if $ab = 1$, the binomial for S is not a recombining tree. If it starts at $S(a, aH)$, and then first goes up and then down it returns to $S(b, abH) = S(b, H) \neq S(a, H)$, but if it first goes down and then up it does return to $S(a, baH) = S(a, H)$.

The state price equations for this system are given by:

$$p_{aa} = \delta f_a a^{-R}$$

$$p_{ab} = \delta(1 - f_a)b^{-R}$$

and

$$p_{ba} = \delta(1 - f_b)a^{-R}$$

$$p_{bb} = \delta f_b b^{-R}$$

These are four independent equations in the four unknowns,

$$\delta, R, f_a, f_b$$

and the solution is given by

$$\begin{pmatrix} f_a \\ f_b \end{pmatrix} = \left(\frac{p_{ab}p_{ba}}{p_{aa}p_{bb}} - 1 \right)^{-1} \begin{pmatrix} \frac{p_{ab}}{p_{bb}} - 1 \\ \frac{p_{ba}}{p_{aa}} - 1 \end{pmatrix}$$

$$R = \frac{\ln \left(\frac{f_a}{1 - f_a} \right) + \ln \left(\frac{p_{aa}}{p_{ab}} \right)}{\ln \left(\frac{b}{a} \right)}$$

and

$$\delta = p_{aa} f_a a^R$$

This example also further clarifies the importance of state dependence. With state independence there are only two equilibrium state equations in the three unknowns, R , f , and δ ,

$$p_a(H) = \delta f a^{-R}$$

$$p_a(H) = \delta(1 - f)b^{-R}$$

Nor can this be augmented by recombining since, assuming $ab = 1$,

$$p_a(bH) = \delta f \left(\frac{1}{b^{-R}} \right) = \delta f b^R = \delta f a^{-R}$$

which is identical to the first equation. In other words, while the parametric assumption has reduced finding the two element kernel to recovering a single parameter, R , it has also eliminated one of the equations. Nor is it possible to recover without the assumption on preferences since while that can allow recombining to provide a third independent equation, it also adds a new variable to the kernel. As we have shown, though, assuming meaningful state dependency once again allows full recovery.

This approach also allows for recovery if the rate of consumption is state dependent. Suppose, for example, that consumption is k_a or k_b in the respective states, a and b . The equilibrium state equations are now

$$p_{aa} = \delta f_a a^{-R}$$

$$p_{ab} = \delta(1 - f_a) \left(\frac{k_a}{k_b}\right)^R b^{-R}$$

$$p_{ba} = \delta(1 - f_b) \left(\frac{k_b}{k_a}\right)^R a^{-R}$$

and

$$p_{bb} = \delta f_b b^{-R}$$

These are four independent equations which can be solved for the four unknowns, δ , R , f_a , and f_b .

To recover the natural distribution in Section 5 below we will use the Recovery Theorem but we could also use the binomial and multinomial recovery models. Which will work better will ultimately be an empirical question.

Section 4: Some Examples, Extensions and Observations

Example 1

Consider a ‘tree’ model with a lognormally distributed payoff at time T and a representative agent with a constant relative risk aversion utility function,

$$U(S_T) = \frac{S_T^{1-R}}{1-R}$$

The future stock payoff, the consumed ‘fruit’ dividend, is lognormal,

$$S_T = e^{(\mu - \frac{1}{2}\sigma^2)T + \sigma\sqrt{T}z} ,$$

where the parameters are as usual and z is a unit standard normal variable.

The pricing kernel is given by

$$\varphi_T = \frac{e^{-\delta T} U'(S_T)}{U'(S)} = e^{-\delta T} \left[\frac{S_T}{S} \right]^{-R}$$

where S is the current stock dividend that must be consumed at time 0.

Given the natural measure and the kernel, state prices are given by

$$P_T(S, S_T) = \varphi_T \left(\frac{S_T}{S} \right) n_T \left(\frac{S_T}{S} \right) = e^{-\delta T} \left[\frac{S_T}{S} \right]^{-R} n \left(\frac{\ln \left(\frac{S_T}{S} \right) - \left(\mu - \frac{1}{2} \sigma^2 \right) T}{\sigma \sqrt{T}} \right)$$

where $n(\cdot)$ is the normal density function.

In this model we know both the natural measure and the state price density and our objective is to see how accurately we can recover the natural measure and, thus, the kernel from the state prices alone using the Recovery Theorem. Setting $T = 1$, Table 9 displays natural transition probability matrix, F, the pricing kernel, the time discount, δ , and the matrix P of transition prices. The units of relative stock movement, S_T/S , are the grid of units of sigma from -5 to +5. Sigma can be chosen as the standard deviation of the derived martingale measure from P, but alternatively we chose the current at-the-money implied volatility from option prices as of March 15, 2011.

With an assumed market return of 8%, a standard deviation of 20% and calculate the characteristic vector of P. As anticipated there is one positive vector and it exactly

equals the pricing kernel shown in Figure 9 and the characteristic root is $e^{-.02} = .9802$ as was assumed. Solving for the natural transition matrix, F , we have exactly recovered the posited lognormal density.

This example fits the assumptions of the Recovery Theorem closely except for having a continuous distribution rather than a discrete one. The closeness of the results with the actual distribution and kernel provides comfort that applying the theorem by truncating the tails is an appropriate approach. Notice that since we can take the truncated portions as the cumulative prices of being in those regions, there is no loss of accuracy in estimating cumulative tail probabilities.

Example 2

The second example deals explicitly with the case where returns are state dependent and can lead to growing values. This would occur, for example, in any model where future prices are some return multiplied by current prices. For this example, once again assume a lognormal distribution but now it is applied to returns and not terminal wealth,

$$S_T = S e^{(\mu - \frac{1}{2}\sigma^2)T + \sigma\sqrt{T}z} ,$$

There are several ways to approach modeling preferences in this example. We could simply assume that there is a representative agent with preferences over consumption which, in this simple example, could be a dividend from the stock price. An easier and more straightforward approach is to assume that preferences are over returns. In the absence of arbitrage this is equivalent in a one period setting to preferences over terminal outcomes. Again using the constant relative risk aversion utility function, the kernel is given by

$$\varphi_T = e^{-\delta T} U' \left(\frac{S_T}{S_0} \right) = e^{-\delta T} \left[\frac{S_T}{S_0} \right]^{-R}$$

While this looks the same as before it is not; here we are not holding the distribution of the terminal value constant, rather we assume a distribution for the relative value.

If the distribution was assumed independent over time, then we would have an awkward problem because the ex ante equilibrium would be the same independent of the return realization and that would make the interest rate constant which would mean that the only equilibrium would be one in which the characteristic vector of implied prices is the same across states, i.e., the model degenerates to risk neutrality.

An easy way to see this is to look closely at the discrete pricing equation

$$U'_i p_{ij} = \delta U'_j f_{ij}$$

Since preferences depend on the relative return the pricing equation is amended to

$$p_{ij} = \delta U'(R_{ij})f_{ij} = \delta U'(R_j)f_j$$

where both the return, R , and the probability are assumed independent of the previous return. It follows, then that

$$p_{ij} = \delta f_j U'(R_j) \equiv p_j$$

independent of state i . Hence, the resulting transition pricing matrix has all the same rows and it must have the same interest rate for all states. This is very similar to the paradox that occurs in an intertemporal equilibrium portfolio model where the opportunity set is unchanging and, as a consequence, asset prices don't change either making any assumed capital gains process inconsistent. Once again, one possible solution is that the natural probabilities equal the martingale probabilities, but the space of possible solutions is $m-1$ dimensional and the natural distribution cannot be recovered.

To make the model interestingly state dependent so that we can apply the Recovery Theorem we make the distribution of current returns depend on past returns through the linkage of returns and implied volatility. A simple regression of implied volatility, v , on past returns produces a very significant relation of the form,

$$\ln(v_{t+1}/v_t) = -\beta \ln(R_t)$$

Where, for this example, we set $\beta = 0.5$. Now the transition probability matrix from state i returns, R_i to state j returns, R_j is given by

$$f_{ij} = f(R_i, R_j) = n \left(\frac{\ln(R_j) - \left(\mu - \frac{1}{2} \sigma^2 R_i^{-2\beta} \right) T}{\sigma R_i^{-\beta} \sqrt{T}} \right)$$

and

$$\varphi = \varphi(R_j) = (R_j)^{-R}$$

which is assumed independent of the initial state, i .

Keep in mind that even though we might have modeled preferences incorrectly, we are working backwards from the observed state prices and have found the unique set of probabilities and prices that could have generated it given the assumed model. Our problem is to interpret the resulting pricing kernel because we it exists by assumption and our construction has found the unique one associated with the observed state prices.

We were able to find the exact solution in this case by applying the Recovery Theorem because we made preferences a function of returns and not of levels as would be more usual in a typical intertemporal consumption model. As such, we didn't need to have the utility function generating the kernel be a constant relative risk aversion function, but, if it is, then the model and the example are identical, i.e., the marginal rates of substitution as a function of levels actually depend only on the relative returns just as if we had assumed that was the case. If we do assume a constant relative risk aversion function, though, we can just proceed as we have and use the Recovery Theorem to extract its single parameter, the constant coefficient of relative risk aversion, R . The numerical results are shown in Table 10.

Example 3

This has the identical return distribution as Example 2, but rather than have preferences have constant relative risk aversion, to make levels matter we assume preferences are represented by a constant absolute utility function. Now the kernel is given by

$$\varphi(S_t, S_{t+1}) = \frac{\delta U'(S_{t+1})}{U'(S_t)} = \frac{-\delta A e^{-AS_{t+1}}}{-A e^{-AS_t}} = \delta e^{-A(S_{t+1}-S_t)} = \delta e^{-AS_t(R_{ij}-1)}$$

where A is the coefficient of absolute risk aversion, S_t is the index value at time t , and R_{ij} is the gross return from the transition from state i to state j . Notice that unlike Example 2, the kernel depends on the level of the index as well as on the return, and the state transition will be a transition from one level to another and one volatility level to another.

For a numerical example we will set $A = 1$ and the initial stock level at 1. We will also truncate the distribution at + 5 and - 5 sigma and assume that the marginal utility is constant beyond those extremes. For empirical work such a truncation is necessary. The results are displayed in Table 11. The recovered kernel and the recovered natural distribution agree perfectly with their assumed values. In a second exercise we changed the marginal utilities at the extreme values, -5 and +5 sigma, and again we were able to perfectly recover the assumed distribution and kernel.

Implicitly in these exercises we are recovering an assumed constant marginal utility beyond the extreme values and we are recovering that value. In practice we will observe the state prices from market prices and we will recover the constant marginal utility that would be consistent with the prices beyond the extremes, e.g., the cumulative distribution of prices above + 5 sigma and below -5 sigma.

Example 4 Binomials – Applying the Recovery Theorem

The natural densities in the above examples are well known limits of binomial processes as the jumps grow more frequent and the jump sizes diminish (see Cox, Ross Rubinstein [1979]). For state dependent binomials we can apply the Multinomial Recovery Theorem, recover the jump probability, f , the subjective discount factor, δ , and the pricing kernel or, more precisely, the projection of the kernel onto the space generated by the binomial. Suppose, though, that the process is state independent. Now we can solve by making an assumption that alters the process by truncating it in the tails for very large and very small outcomes.

Making the appropriate assumption in the tails, it is possible to fully recover the kernel, f and δ , by applying the Recovery Theorem. For simplicity, suppose that the process exactly recovers in k up steps and $n - k$ down steps,

$$a^k b^{n-k} = 1$$

Order the possible states achievable in n or less steps, from $1, \dots, m$ where a given state i is

$$a^t b^v, \quad 0 \leq t + v \leq n$$

Letting $a(i)$ denote the number of up jumps in state i and $b(i)$ denote the number of down jumps, for all but the bottom and top states, $(i,j) \neq (1,1)$ or (m,m) , the state probability transition matrix is given by

$$\begin{aligned} f_{ij} &= f \text{ if } (a(j) = a(i) + 1, b(j) = b(i)), \\ &1 - f \text{ if } (a(j) = a(i), b(j) = b(i) + 1) \\ &\text{and } 0 \text{ otherwise} \end{aligned}$$

For the bottom state, we set

$$f_{11} = 1 - f$$

and for the top state we set

$$f_{mm} = f$$

In other words, we truncate the process at 1 and m and lump together the infinity of states below 1 and above m .

Since the process is recurrent in n steps the matrix F is irreducible and we can now fully recover F and, therefore, f , δ , and the pricing kernel, \tilde{Q} . Given that we were unable to recover the process from the state equations for the simplest recombining process, this is a quite surprising result. Let us look again, then, at the simple recombining example.

For this example, let's truncate the process after a single step, i.e., lump all the states above aS and all below bS. Since $ab = 1$, any state is of the form

$$a^i b^j = a^{i-j} \text{ if } i > j \quad b^{j-i} \text{ if } i < j \text{ and } 1 \text{ if } i = j$$

Thus, there are no states between a and b other than $S = 1$, the starting state and state 1 is b, state 2 is 1, and state 3 is S.

The transition probability matrix is given by

$$F = \begin{bmatrix} 1-f & f & 0 \\ 1-f & 0 & f \\ 0 & 1-f & f \end{bmatrix}$$

The projection of the kernel onto the three states is given by

$$\varphi = \begin{pmatrix} U'_1 \\ U'_2 \\ U'_3 \end{pmatrix} = \begin{pmatrix} 1 \\ \varphi_2 \\ \varphi_3 \end{pmatrix}$$

where, without loss of generality, we scale the first element to be unity. Defining D as the diagonal matrix with the kernel elements on the diagonal, the pricing matrix, P, is given by

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

and since F is irreducible, P is as well and we can apply the Recovery Theorem to obtain the underlying parameters. But, how is this possible given our previous results for this simple model requiring state dependence or, to put it somewhat differently what equation did we add to close the system?

Ignoring zeros, we have the following 6 state equations in the original 4 unknowns, f , δ , φ_2 and φ_3 .

$$p_{11} = \delta(1-f)$$

$$p_{12} = \delta\varphi_2 f$$

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

$$\varphi_2 p_{23} = \delta\varphi_3 f$$

$$\varphi_3 p_{32} = \delta\varphi_2(1-f)$$

$$p_{33} = \delta f$$

This system can be solved by applying the Recovery Theorem or, equivalently by solving it directly. The solution is

$$\delta = p_{11} + p_{33}$$

$$f = \frac{p_{33}}{p_{11} + p_{33}}$$

$$\varphi_2 = \frac{p_{11}}{p_{21}}$$

and

$$\varphi_3 = \frac{p_{11}^2}{p_{21}p_{32}}$$

The Recovery Theorem was able to solve the system because we closed the model and put it into an irreducible form. The third and fourth equations govern the transition from the current state, $S = 1$, up and down and they are identical to the two original state equations for the first jump. Combining the second and third equations or, equivalently, the fourth and fifth equations, is identical to the new equation introduced by the recombination in the original equation system. The new equations that are added are the first and the last which lump together the states at or below b and at or above a respectively, and it is the addition of these equations which enables us to solve the system. In other words, the market is assumed to be indifferent between going from the really calamitous to the even more calamitous and between nirvana and more nirvana, fixing the marginal rates of substitution in the extreme mega states unity.

Alternatively, we could set the (undiscounted) kernel at a for the lower mega state and b for the upper mega state. The equilibrium state equations are now

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

where A is a diagonal matrix with $a_{11} = a$, $b_{mm} = b$, and $a_{ii} = 1$, $i \neq 1, m$. Applying the Recovery Theorem to PA^{-1} we can recover the kernel, F and δ as before. A natural way to determine A would be to make estimate a and b as functions of the interior kernel elements, i.e., $A = A(D)$. The problem would then be a fixed point one of finding D that satisfies the above state equations, e.g., for a given $D(j)$ we determine $A(D)$, apply the Recovery Theorem, find $D(j+1)$, and iterate until the process converges.

Section 5: The Risk Neutral Transition Pricing Matrix

With the rich market for derivatives on the S&P 500 index, the market is effectively complete along dimensions related to the index, both value and the states of the return process. The Recovery Theorem relies on knowledge of the martingale transition matrix and in this section we will show how to obtain this information from option prices. Given the wide spread interest in estimating models for pricing derivative securities it is not surprising that the literature on this problem is very extensive (see Rubinstein [1994], Rubinstein and Jackwerth [1996] and Jackwerth [1999], Derman and Kani [1994] and [1998], Dupire [1994]). We draw on only the most basic findings of this work and make no effort to extend it.

We first check that option prices are, in fact, arbitrage free. We begin by clearing up some loose ends on no arbitrage (NA) in the options market.

The Vol Surface and NA

Figure 9 displays the surface of implied volatilities on S&P puts and calls, the 'vol surface', on March 20, 201 drawn as a function of time to maturity, 'tenor', and the strike. Option prices are typically quoted in terms of implied volatilities from the Black-Scholes, i.e., the volatilities that when put into the model give the market premium for the option, but doing so in no way is a statement of the validity of the Black-Scholes model, rather it is simply a transformation of the market determined premiums. The source of the data used in this paper is a bank over the counter bid/offer sheet. While the data is in broad agreement with exchange traded options, we chose this source since the volume on the over the counter market is approximately six times that on the exchange even for common at-the-money contracts.

The surface displays a number of familiar features. There is a 'smile' with out of the money and in the money options having the highest implied volatilities. The shape is actually a 'smirk' with more of a rise in implied volatility for out of the money puts (in the money calls). One explanation for this is that there is an excess demand for out of the money puts to protect long equity positions relative to the expectations the market has about future volatilities. Notice, too, that the surface has the most pronounced curvature for short dated options and that it rises and flattens out as the tenor increases. A story supporting this is the demand for long dated calls by insurance companies that have sold variable annuities. Whatever the merit of these explanations, these are persistent features of the vol surface at least since the crash in 1987. As an empirical matter, as we employed in Example 2 in the previous section, implied vols move inversely with the market, and, more generally, the surface changes stochastically over time, but the basic shape displayed in Figure 12 is common.

Implied volatilities are a function of the risk neutral probabilities, the product of the natural probabilities and the pricing kernel - risk aversion. Does the skew indicate that the market assigns a high probability to a significant decline or does it mean that the probability of that occurring is no higher than the long run average of frequency of such

events, but that the risk aversion is very high? We will apply the Recovery Theorem to the state prices implicit in the vol surface to answer this question but first we clean some loose ends in the literature on option prices.

No Arbitrage (NA) and Implied Vols

To use the information in the vol surface we first must check that it doesn't admit of arbitrage. It is well known (see Merton [1974]) that for any given tenor the vol surface has to be downward sloping and convex as a function of the strike. The set of necessary conditions for options (of a given tenor) to be arbitrage free is:

The Option Conditions (OC)

Letting $C(K)$ denote the price of a call option as a function of its strike, the familiar necessary conditions for the absence of arbitrage are:

- $C(K) > 0$
- $C'(K) < 0$
- $C''(K) > 0$
- $C(K) \downarrow 0$ as $K \rightarrow \infty$
- $C(0) = S$ (the initial stock price)

Theorem 5

No Arbitrage (NA) if and only if the OC are satisfied.

Proof:

Necessity is well known (see Merton (1974)). For sufficiency, let α denote a proposed arbitrage portfolio satisfying:

$$\int_0^s [s - x]\alpha(x)dx \equiv g(s) > 0, a. e.$$

with strict inequality on a set of positive measure (we are only considering portfolios with adequate regularity conditions and not considering portfolios with atoms). From Breeden Litzenberger (1978), $\alpha(s) = g''(s)$. If $C(k)$ denotes the cost of a call with a fixed maturity and with a strike of k , it follows that the cost of the portfolio α is:

$$\int_0^\infty C(k)\alpha(k)dk = \int_0^\infty C(k)g''(k)dk = C(\infty)g'(\infty) - C(0)g'(0) - \int_0^\infty C'(k)g'(k)dk$$

Since

$$g'(k) = \int_0^k \alpha(x)dx \text{ hence } g'(0) = 0$$

Restricting our attention to α such that $C(K)g'(K) \rightarrow 0$ as $K \uparrow$ we have

$$\begin{aligned}\int_0^\infty C(k)\alpha(k)dk &= \int_0^\infty C(k)g''(k)dk \\ &= -\int_0^\infty C'(k)g'(k)dk = -C'(\infty)g(\infty) + C'(0)g(0) + \int_0^\infty C''(k)g(k)dk\end{aligned}$$

Denoting the state price density by $p(S)$ we have

$$C(K) = \int_K^\infty (S - K)p(S)dS \text{ hence } C'(K) = -\int_K^\infty p(S)dS$$

And

$$C'(0) = -\int_0^\infty p(S)dS = -e^{-rT} \text{ and } C'(\infty) = 0$$

By assumption $g > 0$ and since $C''(K) > 0$ it follows that the portfolio must have a positive cost making arbitrage impossible.

□

For any given surface it is straight forward to check that the OC conditions are satisfied.

Since all contracts can be formed as portfolios of options (Ross [1976]) it is not surprising that from the vol surface we can derive the state price distribution, $p(S)$:

$$C(K) = \int_0^\infty [S - K]^+ p(S)dS = \int_K^\infty [S - K]p(S)dS$$

and differentiating twice we obtain the Breeden Litzenberger [1978]) result that

$$p(S) = C''(K)$$

Numerically approximating this second derivative as a second difference along the surface at each tenor yields the distribution of state prices looking forward from the current state. Setting the grid size of index movements at 0.5% the S&P 500 call options on April 27, 2011 produced the marginal state price densities reported in Table 12. The results are broadly sensible with the exception of the relatively high implied interest rates at longer maturities which we will address below.

Section 6: The Recovered Natural Distribution,

Unfortunately, there isn't a rich forward market for options and we don't directly observe the price transition function, $P = [p_{ij}]$. Instead we will derive it from the state price distributions at different tenors.

$$Q(x, y, T) = \int Q(x, z, t)Q(z, y, T - t)$$

where $Q(x, y, \tau)$ is the forward probability transition function for going from state x to state y in τ periods and where the integration is over the intermediate state at time t, z .

The forward equation for state prices is given by:

$$p(x, y, T) = \int p(x, z, \tau)P(z, y, T - \tau)$$

where $p(x, y, T)$ is the state price of one dollar in state y at time T given that at time 0 the state is x , and $P(z, y, T - t)$ is the transition price function of a dollar in state y given that the current state is x and $T - t$ is the remaining time. With appropriate regularity conditions this can be solved for P given that we know the current p vectors. For a discrete system with m states we have the $m \times m$ matrix:

$$P = [p(i, j)] \text{ denotes pricing from state } i \text{ to state } j$$

At the current date we are in some state, θ , and observe the current prices of options across strikes and tenors, and, as shown above, we can extract from the state prices at each future date T ,

$$p^T = \langle p(1, T), \dots, p(m, T) \rangle$$

Letting the stock price and the past return index the states (this allows, for example, for states to depend on price paths) and denoting the current stock price as S , p^1 is column S of P . To solve for the remaining elements of P we apply the forward equation recursively to create the sequence of $m-1$ equations:

$$p^{t+1} = p^t P, t = 1, \dots, m$$

This is a system of m^2 individual equations in the m^2 variables P_{ij} and since we know the current prices, p^t it can be solved by recursion.

Since our intention is illustrative we ignore the potential state dependence on past returns and identify the states only by the price level. For relatively short periods this may not be much different than if we also used returns since the final price over, say a quarter, is a good surrogate for the price path. Friday, April 27, 2011. The grid is chosen to be from -5 to $+5$ standard deviations with a standard deviation of 9%/quarter. This

seemed a reasonable compromise between fineness and coverage in the tails. The analysis above was then implemented numerically to derive the transition pricing matrix, P by varying the choice of P so as to minimize the sum of squared deviations between the resulting prices and the state price vectors of Table 12. The resulting forward transition price matrix, P , is shown in Table 12 under the table of the marginal state price densities.

The final step is to apply the Recovery Theorem to the transition pricing matrix, P . Table 13 displays the recovered pricing kernel and the resulting natural probability transition matrix. The kernel declines monotonically as the stock value rises, but this need not be the case. Interestingly the nearly block diagonal form for the probability transition matrix mirrors the subordinated log normal with volatility dependent on past returns of Examples 2 and 3 of Section 4. Figure 10 displays the pricing kernel, normed with the current marginal utility set at 1, and compares it with a best fit constant relative risk aversion utility function chosen to minimize the sum of squared deviations. Notice that the fitted kernel function is higher in the region of low and high returns than the recovered kernel indicating that its impact on pricing will be exaggerated in these regions. There are no doubt better utility functions to use, but it does give some pause about the ubiquitous use of these functions in intertemporal financial portfolio analysis. Intriguingly, the best fit is close to a log utility function, i.e., the growth optimal criterion.

Table 14 shows the recovered marginal distributions at the future dates and compares them with the historical distribution estimated by a bootstrap of S&P 500 returns from 60 years of data (1960 – 2010). Table 14 also displays the implied volatilities from the option prices on April 27, 2011. Not surprisingly, the densities display a slight upward drift as the horizon (tenor) increases. The summary statistics display some significant differences between the recovered and the historical distributions. For the recovered, which is a forward looking measure, the annual expected return at all horizons is approximately 5%/year as compared with 10%/year for the historical measure. The standard deviation, on the other hand, is comparable at about 15%/ year; a not surprising result given the greater accuracy inherent in implied volatilities and the fact that with diffusions they coincide – albeit with bias - more closely with realized volatilities than do expected returns and realized returns.

The state prices should sum to the riskless interest factor. The rates are relatively accurate out to about 1 but then rise from 1.85% at 1 year to 7.93% at 3 years. This is significantly higher than 3 year (swap) rates at the time and indicative of a bias in the computation of the state prices. With this high and incorrect rate the risk premium turns negative 2 years out, and, the Sharpe ratio also must turn negative as well. Some of this is explainable by the omission of dividends, but it is no doubt mostly a consequence of the error in our computation of state prices and the risk free rates and speaks to the need for a more refined estimation.

Notice that the at-the-money implied volatilities are significantly higher than those derived from the recovered distribution. This is a phenomena closely related to the observation that implied volatilities are generally significantly greater than realized

volatility and it is not surprising that the volatilities from the recovered distribution have a similar relation to realized volatility.

Table 15 displays and compares the recovered natural density and distributions with those obtained from the bootstrap, and Figure 11 plots these densities. Of particular interest is what they say about the long standing concern with tail events. Rietz [1988] argued that a large but unobserved probability of a catastrophe – ‘tail risk’ - could explain the equity risk premium puzzle, i.e., the apparent dominance of stocks over bonds and related questions. Barro [2006] lent support to this view by expanding the data set to include a wide collection of catastrophic market drops beyond what one would see with a single market and Weitzmann [2007] provided a deep theoretical argument in support of fat tails. Somewhat more pithily, Merton Miller observed after the 1987 crash that 10 standard deviation events seemed to be happening every few years.

As we said in the introduction, tail risk is the economists’ version of the cosmologists’ dark matter. It is unseen and not directly observable but it exerts a force that can change over time and that can profoundly influence markets (or galaxies). By separating the risk-averse kernel from the subjective probabilities embedded in option prices, though, we get can shed some light on the dark matter and estimate the market’s probability of a catastrophe. As Figure 11 show, the recovered density has a fatter left tail than the historical distribution. Table 15 puts the probability of a six month drop in excess of 32% at .0008 or 4 in 5000 bootstraps. By contrast, the recovered density puts this probability at 1.2%. Similarly, the historical probability of a drop in excess of 26% in a six month period is .002 (10 times in 5000 bootstraps) while the recovered market probability of .0223 is 10 times greater at over 2%.

This is only a crude first pass at applying the Recovery Theorem, and it is intended to be indicative rather than conclusive. There is an enormous amount of work to be done beginning with a more careful job of estimating the state price density from option prices and, from there, estimating the state price transition matrix from the state price density at different horizons and strikes. There are also many improvements required to accurately recover the kernel and the natural measure implicit in the state prices.

Section 7: Testing the Efficient Market Hypothesis

It has long been thought that tests of efficient market hypotheses are necessarily joint tests of both market efficiency and a particular asset pricing model. Using the Recovery Theorem we can separate these two assumptions and derive model free tests of the efficient market hypothesis.

One way of approaching this question actually requires no such separation. A somewhat traditional approach makes use of the data in option prices directly. Let $x(S)$ be the (date T) payoff on a proposed strategy for ‘beating the market’. By assumption, $x(S)$ is a self-financing strategy so the null hypothesis is that the market is efficient and that any self-financing strategy must have zero current value:

$$E^*[x] = \int_0^\infty x(S)p(S)dS = 0$$

But, any strategy $x(S)$ can be written as a combination of call options

$$x(S) = \mu + \int_0^\infty [S - x]^+ \beta(x) dx + \epsilon$$

where the β coefficients are from the projection of $x(S)$ on the calls. The value of x is now given by

$$\begin{aligned} 0 = E^*[x] &= e^{-rT} \mu \\ &+ \int_0^\infty \left[\int_0^\infty [z - K]^+ \beta(K) dK \right] \varphi(z) dz \\ &= \int_0^\infty \left[\int_0^\infty [z - x]^+ \varphi(z) dz \right] \beta(K) dK = e^{-rT} \mu + \int_0^\infty \beta(K) C(K) dK \end{aligned}$$

The test of efficiency is simply a test of the above restriction

$$e^{-rT} \hat{\mu} + \int_0^\infty \hat{\beta}(K) C(K) dK = 0$$

More interestingly, though, we can follow through on a completely different tack. In Ross [2005] an approach to testing efficient market hypotheses was proposed that depended on finding an upper bound to the volatility of the pricing kernel.

Assume that μ is stochastic and depends on some unspecified or unobserved conditioning information set, I . From the Hansen – Jagannathan bound [1991] we have a lower bound on the volatility of the pricing kernel

$$\sigma(\varphi) \geq (e^{-rT}) \frac{\mu}{\sigma}$$

where μ is the excess return and σ is the standard deviation on any asset which implies that $\sigma(\varphi)$ is bounded from below by the discounted maximum observed Sharpe ratio.

Equivalently, this is also an upper bound on the Sharpe ratio for any investment. From the recovered marginal density function reported in Table 14 we can compute the variance of the kernel at, say, one year out. The computation is straightforward and the resulting variance is

$$\sigma^2(\varphi) = .1065$$

or an annual standard deviation of

$$\sigma(\varphi) = .3264$$

Ignoring the small interest factor, this is the upper limit for the Sharpe ratio for any strategy to be consistent with efficient markets and not be ‘too good’ a deal (see Cochrane [1999] and Bernado and Ledoit [1999] for a discussion of good deals and Ross [1976] for an early use of the bound for asset pricing).

Similarly, we can decompose excess returns on some asset or portfolio strategy as

$$x_t = \mu(I_t) + \epsilon_t ,$$

where the mean depends on the particular information set, I, and where the residual term

$$\sigma^2(x_t) = var(\mu(I_t)) + \sigma^2(\epsilon_t) \leq E[\mu^2(I_t)] + \sigma^2(\epsilon_t)$$

Hence, we have an upper bound to the R^2 of the regression in an efficient market

$$R^2 = \frac{\sigma^2(\mu(I_t))}{\sigma^2(x_t)} \leq \frac{E[\mu^2(I_t)]}{\sigma^2(x_t)} \leq e^{2rT} \sigma^2(\varphi)$$

i.e., the R^2 is bounded above by the volatility of the pricing kernel. Of course, the kernel can have arbitrarily high volatility by simply adding orthogonal noise to it, so the proper maximum to be used is the volatility of the projection of the kernel on the stock market, and, hence, these are tests on strategies that are based on stock returns and the filtration they generate.

Using our estimate of the variance of the pricing kernel we find that the maximum it can contribute to the R^2 of an explanatory regression is about 10%. In other words, 10% of the annual variability of an asset return is the maximum amount that can be attributed to movements in the pricing kernel and 90% should be idiosyncratic in an efficient market. (This bound is the same for any time unit of observation.) Hence any test of an investment strategy that uses, say, publicly available data, and has the ability to predict future returns with an $R^2 > 10\%$ would be a violation of efficient markets independent of the asset pricing model being used. Of course, any such strategy must overcome transactions cost to be an implementable violation, and a strategy that could not overcome those costs would be purely of academic interest.

Section 8: Summary and Conclusions

If we can observe or estimate the transition price matrix then the Recovery Theorems allow us to recover the pricing kernel, i.e., the marginal rates of substitution of the representative agent, and the (ex ante) natural distribution of asset returns implicit in market prices. Armed with the market's risk aversion – the pricing kernel – and the market's subjective assessment of the distribution of returns, there is a cornucopia of applications. Currently economists and investors are regularly asked to fill out surveys to determine some consensus estimate for the expected return on the stock market. Now we can directly assess this much as we use forward rates as forecasts of future spot rates. Institutional asset holders, such as pension funds, use historical estimates of the risk premium on the market as an input into asset allocation models. The market's current subjective forecast would probably be superior and certainly of interest. Risk control models such as VAR, typically use historical estimates to determine the risk of various books of business and this, too, would be improved by using the recovered distribution.

These results are also applicable across a variety of markets, e.g., currency, futures, and fixed income. For the stock market, they can be used to examine the host of market anomalies and, more specifically, market efficiency for which we have presented one of many possible tests based on the recovered distribution. The ability to better assess the market's perspective of the likelihood of a catastrophic drop will have both practical and theoretical implications. The kernel is important on its own since it measures the degree of risk aversion in the market. For example, just as the market portfolio is a benchmark for performance measurement and portfolio selection, the pricing kernel serves as a benchmark for preferences. Knowledge of both the kernel and the natural distribution will also shed light on the controversy of whether the market is too volatile to be consistent with rational pricing models (see, e.g., Leroy and Porter [1981], Shiller [1981]).

There is certainly much work to be done on the theoretical front. We need to further explore recovery along binomial (multinomial) trees. We don't yet know what approach will prove the most practical and applicable version of the Recovery Theorem. We also need to examine the recovery theory in a continuous time and state setting and, in particular, we need to specialize it to particular processes such as state dependent diffusions where recovery will be about the parameters governing the return processes. More generally, we need to carefully examine the boundaries of recovery, and particularly the robustness of the underlying assumptions on the existence of the state independent kernel.

In conclusion, contrary to finance folklore, it is possible to separate risk aversion from the natural distribution and estimate each of them from market prices. With a pun intended, we have only scratched the surface of discovering what investors are forecasting for the future development of the market and the economy.

Figure 1

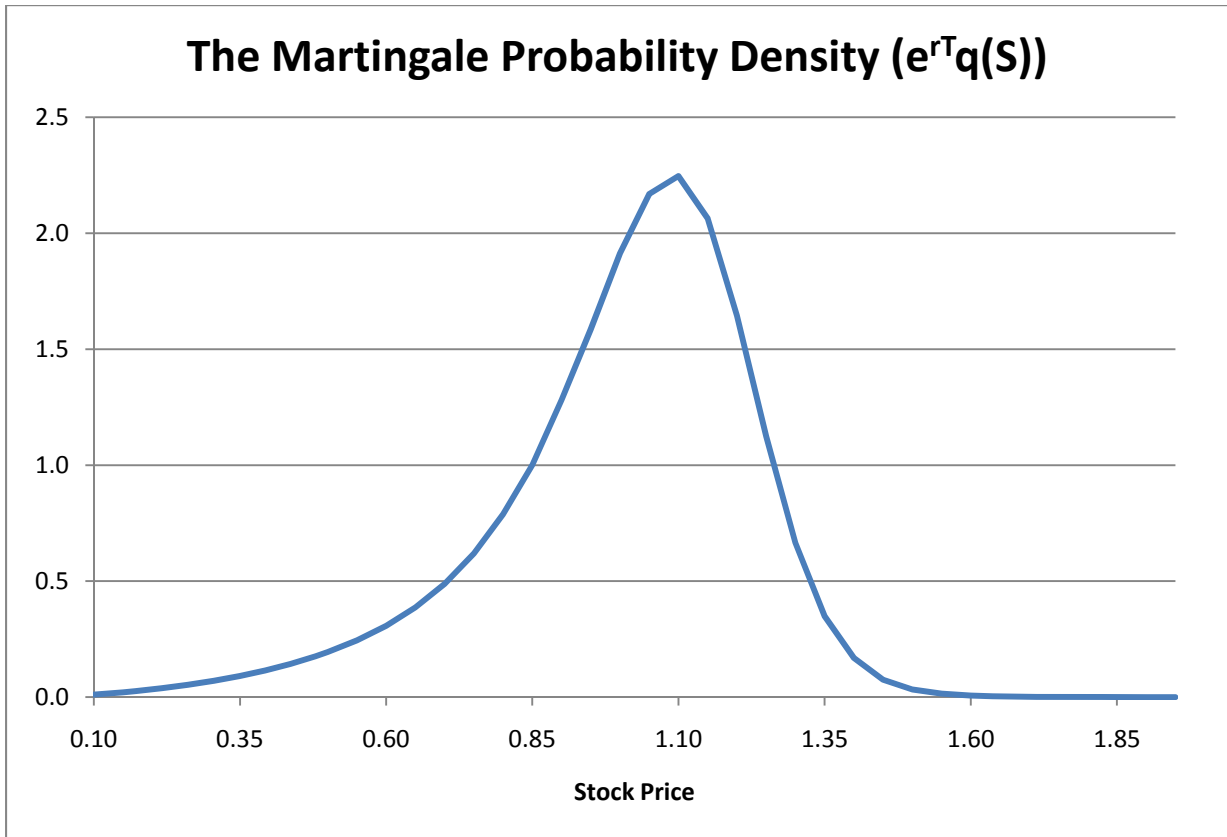


Figure 2




Figure 3

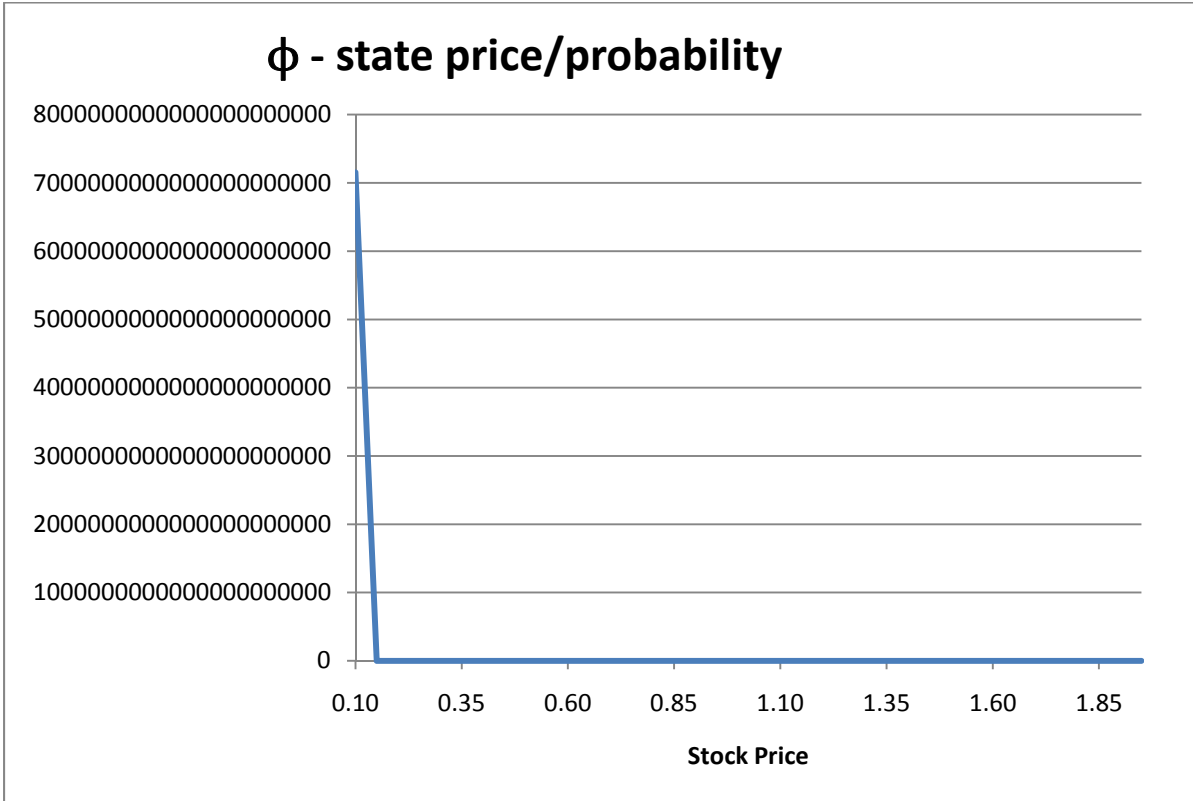


Figure 4

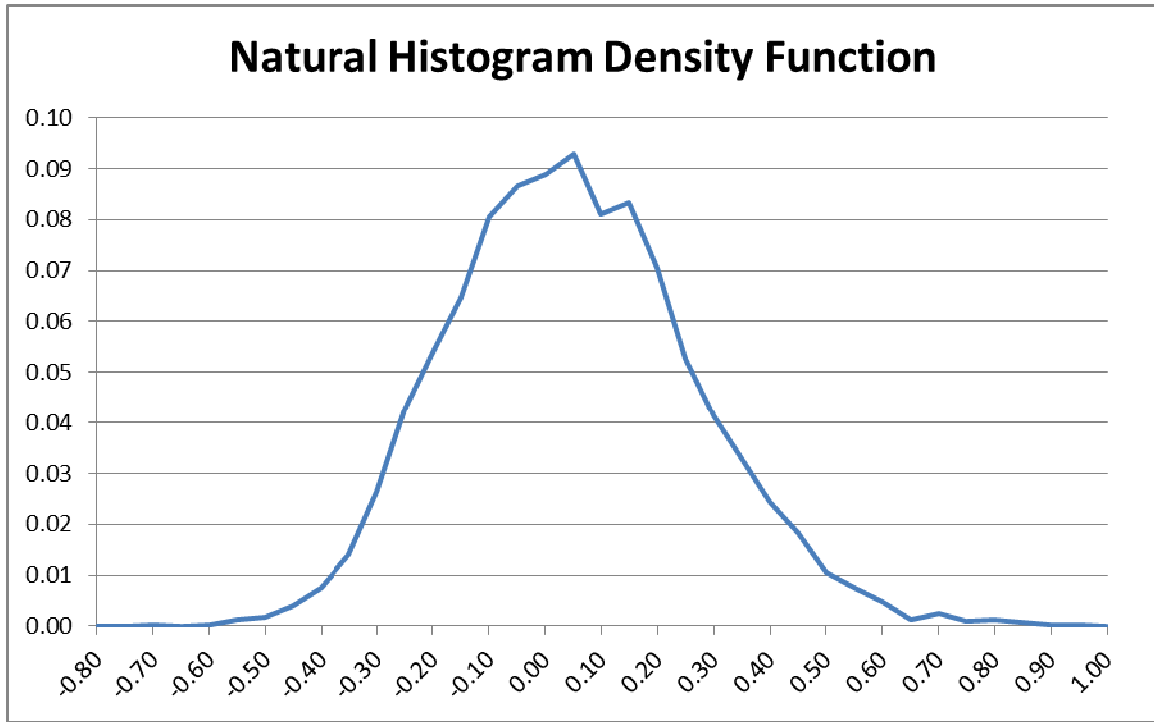


Figure 5




Figure 6

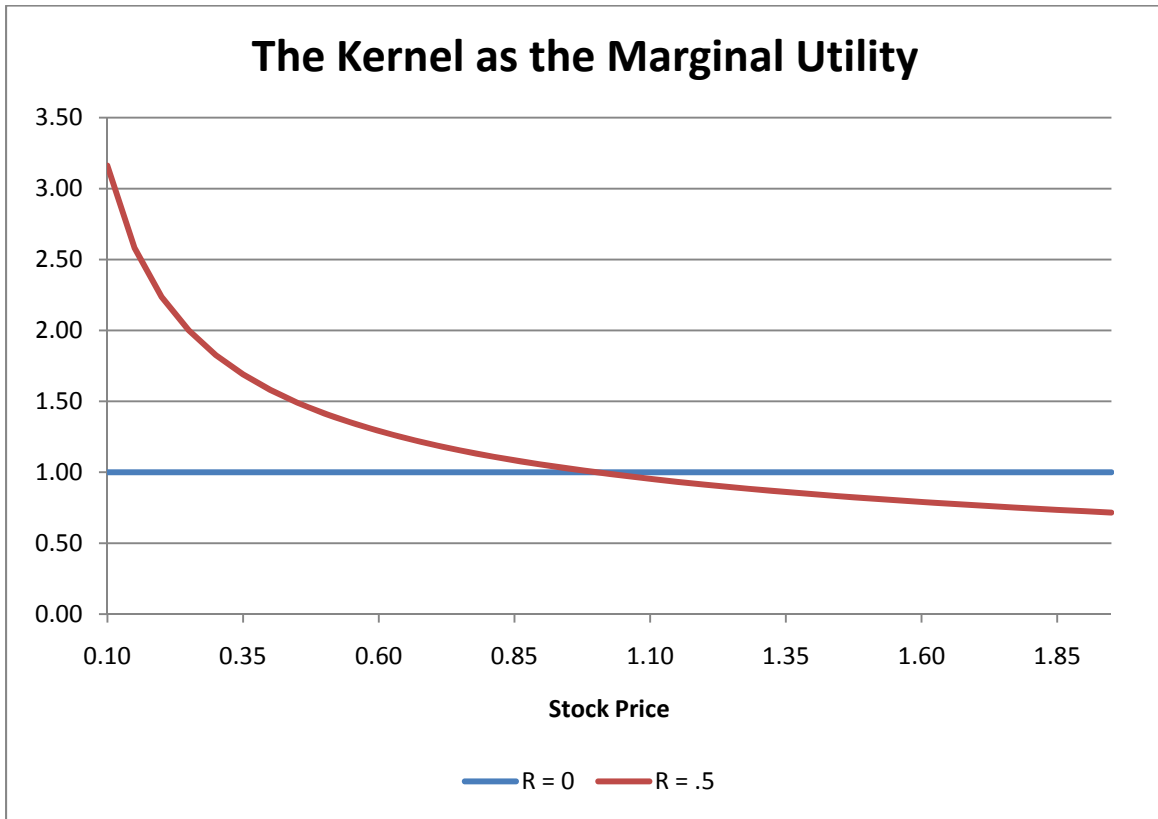


Figure 7

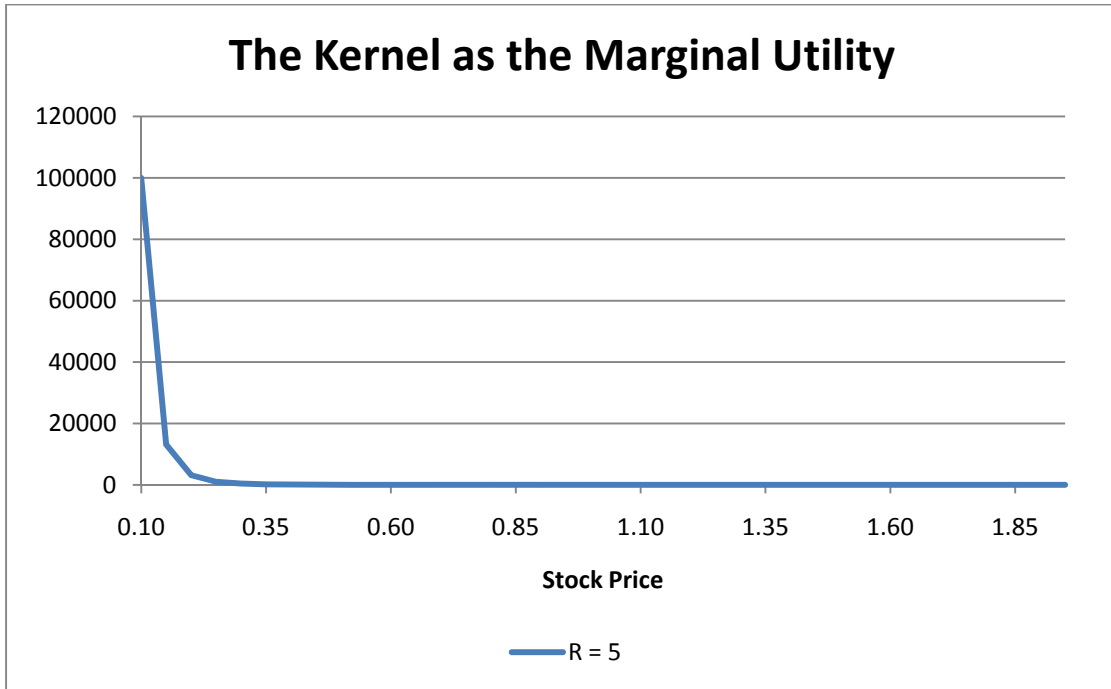


Figure 8

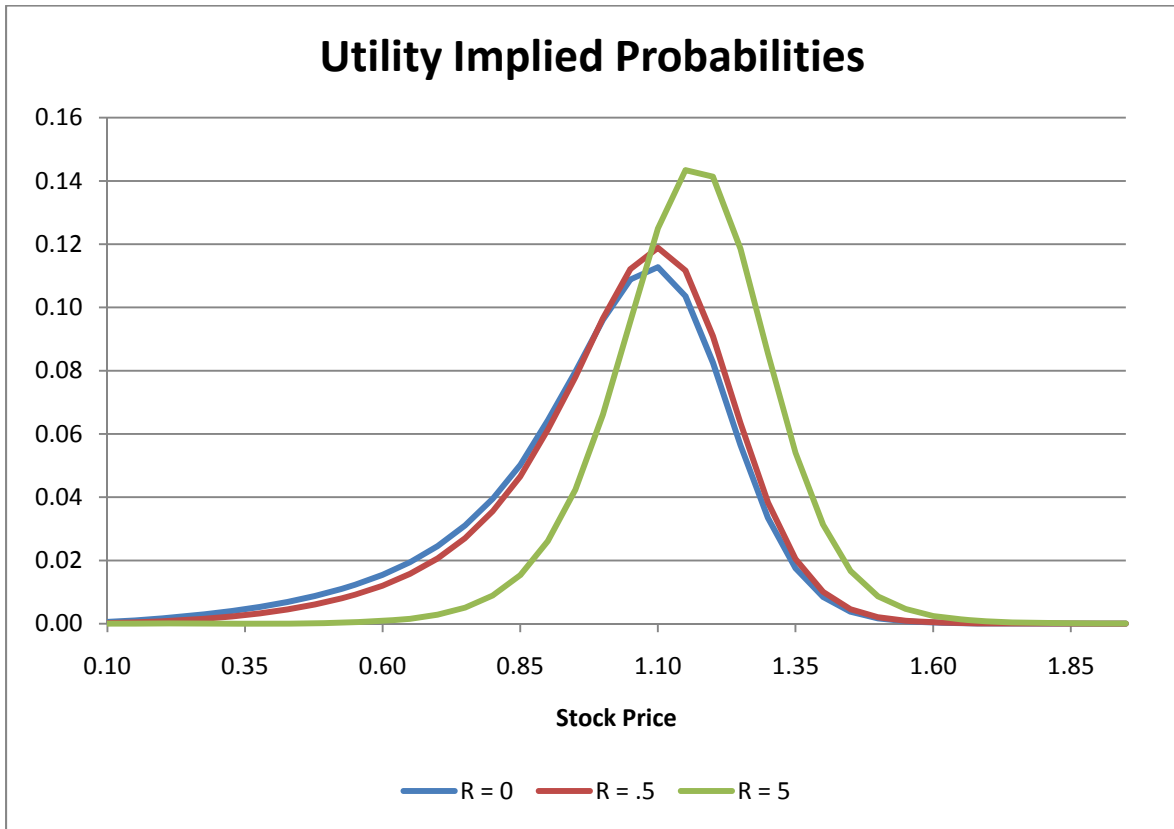


Table 9

Fixed future payoff - Lognormal Distribution, Constant Relative Risk Aversion, R = 3												
The State Space Transition Matrix, P												
	Sigmas	-5	-4	-3	-2	-1	0	1	2	3	4	5
Sigmas	$S_0 \setminus S_T$	0.37	0.45	0.55	0.67	0.82	1.00	1.22	1.49	1.82	2.23	2.72
-5	0.37	0.0000	0.0000	0.0005	0.0046	0.0152	0.0186	0.0084	0.0014	0.0001	0.0000	0.0000
-4	0.45	0.0000	0.0000	0.0009	0.0084	0.0278	0.0339	0.0152	0.0025	0.0002	0.0000	0.0000
-3	0.55	0.0000	0.0001	0.0017	0.0152	0.0506	0.0618	0.0278	0.0046	0.0003	0.0000	0.0000
-2	0.67	0.0000	0.0001	0.0031	0.0278	0.0922	0.1126	0.0506	0.0084	0.0005	0.0000	0.0000
-1	0.82	0.0000	0.0002	0.0056	0.0506	0.1680	0.2052	0.0922	0.0152	0.0009	0.0000	0.0000
0	1.00	0.0000	0.0004	0.0102	0.0922	0.3061	0.3738	0.1680	0.0278	0.0017	0.0000	0.0000
1	1.22	0.0000	0.0008	0.0186	0.1680	0.5577	0.6812	0.3061	0.0506	0.0031	0.0001	0.0000
2	1.49	0.0000	0.0014	0.0339	0.3061	1.0162	1.2412	0.5577	0.0922	0.0056	0.0001	0.0000
3	1.82	0.0000	0.0025	0.0618	0.5577	1.8516	2.2616	1.0162	0.1680	0.0102	0.0002	0.0000
4	2.23	0.0001	0.0046	0.1126	1.0162	3.3739	4.1209	1.8516	0.3061	0.0186	0.0004	0.0000
5	2.72	0.0001	0.0084	0.2052	1.8516	6.1476	7.5087	3.3739	0.5577	0.0339	0.0008	0.0000
$\pi =$		20.09	11.02	6.05	3.32	1.82	1.00	0.55	0.30	0.17	0.09	0.05
The Natural Probability Transition Matrix, F												
	Sigmas	-5	-4	-3	-2	-1	0	1	2	3	4	5
Sigmas	$S_0 \setminus S_T$	0.37	0.45	0.55	0.67	0.82	1.00	1.22	1.49	1.82	2.23	2.72
-5	0.37	0.0000	0.0000	0.0017	0.0283	0.1714	0.3814	0.3123	0.0940	0.0104	0.0004	0.0000
-4	0.45	0.0000	0.0000	0.0017	0.0283	0.1714	0.3814	0.3123	0.0940	0.0104	0.0004	0.0000
-3	0.55	0.0000	0.0000	0.0017	0.0283	0.1714	0.3814	0.3123	0.0940	0.0104	0.0004	0.0000
-2	0.67	0.0000	0.0000	0.0017	0.0283	0.1714	0.3814	0.3123	0.0940	0.0104	0.0004	0.0000
-1	0.82	0.0000	0.0000	0.0017	0.0283	0.1714	0.3814	0.3123	0.0940	0.0104	0.0004	0.0000
0	1.00	0.0000	0.0000	0.0017	0.0283	0.1714	0.3814	0.3123	0.0940	0.0104	0.0004	0.0000
1	1.22	0.0000	0.0000	0.0017	0.0283	0.1714	0.3814	0.3123	0.0940	0.0104	0.0004	0.0000
2	1.49	0.0000	0.0000	0.0017	0.0283	0.1714	0.3814	0.3123	0.0940	0.0104	0.0004	0.0000
3	1.82	0.0000	0.0000	0.0017	0.0283	0.1714	0.3814	0.3123	0.0940	0.0104	0.0004	0.0000
4	2.23	0.0000	0.0000	0.0017	0.0283	0.1714	0.3814	0.3123	0.0940	0.0104	0.0004	0.0000
5	2.72	0.0000	0.0000	0.0017	0.0283	0.1714	0.3814	0.3123	0.0940	0.0104	0.0004	0.0000

Table 10

Returns - Lognormal Distribution, Constant Relative Risk Aversion of returns, R = 2												
The State Space Transition Matrix, P												
	Sigmas	-5	-4	-3	-2	-1	0	1	2	3	4	5
Sigmas	$S_0 \setminus S_T$	0.37	0.45	0.55	0.67	0.82	1.00	1.22	1.49	1.82	2.23	2.72
-5	0.37	0.0000	0.0001	0.0015	0.0136	0.0442	0.0528	0.0232	0.0037	0.0002	0.0000	0.0000
-4	0.45	0.0000	0.0001	0.0020	0.0186	0.0630	0.0784	0.0359	0.0060	0.0004	0.0000	0.0000
-3	0.55	0.0000	0.0001	0.0026	0.0254	0.0895	0.1163	0.0555	0.0098	0.0006	0.0000	0.0000
-2	0.67	0.0000	0.0001	0.0034	0.0344	0.1268	0.1720	0.0859	0.0158	0.0011	0.0000	0.0000
-1	0.82	0.0000	0.0002	0.0044	0.0463	0.1787	0.2540	0.1328	0.0255	0.0018	0.0000	0.0000
0	1.00	0.0000	0.0002	0.0056	0.0618	0.2506	0.3738	0.2052	0.0414	0.0031	0.0001	0.0000
1	1.22	0.0000	0.0002	0.0070	0.0818	0.3491	0.5484	0.3169	0.0674	0.0053	0.0002	0.0000
2	1.49	0.0000	0.0003	0.0087	0.1070	0.4827	0.8012	0.4892	0.1099	0.0091	0.0003	0.0000
3	1.82	0.0000	0.0003	0.0106	0.1382	0.6615	1.1648	0.7546	0.1798	0.0158	0.0005	0.0000
4	2.23	0.0000	0.0003	0.0127	0.1758	0.8968	1.6833	1.1624	0.2953	0.0276	0.0009	0.0000
5	2.72	0.0000	0.0004	0.0148	0.2196	1.2007	2.4151	1.7871	0.4865	0.0487	0.0018	0.0000
$\bar{\sigma} =$		7.39	4.95	3.32	2.23	1.49	1.00	0.67	0.45	0.30	0.20	0.14
The Natural Probability Transition Matrix, F												
	Sigmas	-5	-4	-3	-2	-1	0	1	2	3	4	5
Sigmas	$S_0 \setminus S_T$	0.37	0.45	0.55	0.67	0.82	1.00	1.22	1.49	1.82	2.23	2.72
-5	0.37	0.0000	0.0001	0.0035	0.0461	0.2232	0.3977	0.2607	0.0629	0.0056	0.0002	0.0000
-4	0.45	0.0000	0.0001	0.0031	0.0423	0.2133	0.3961	0.2706	0.0680	0.0063	0.0002	0.0000
-3	0.55	0.0000	0.0001	0.0027	0.0386	0.2033	0.3938	0.2807	0.0736	0.0071	0.0003	0.0000
-2	0.67	0.0000	0.0001	0.0023	0.0351	0.1929	0.3906	0.2909	0.0797	0.0080	0.0003	0.0000
-1	0.82	0.0000	0.0000	0.0020	0.0316	0.1823	0.3865	0.3015	0.0865	0.0091	0.0004	0.0000
0	1.00	0.0000	0.0000	0.0017	0.0283	0.1714	0.3814	0.3123	0.0940	0.0104	0.0004	0.0000
1	1.22	0.0000	0.0000	0.0015	0.0251	0.1600	0.3750	0.3233	0.1025	0.0120	0.0005	0.0000
2	1.49	0.0000	0.0000	0.0012	0.0220	0.1483	0.3673	0.3346	0.1121	0.0138	0.0006	0.0000
3	1.82	0.0000	0.0000	0.0010	0.0191	0.1362	0.3579	0.3459	0.1230	0.0161	0.0008	0.0000
4	2.23	0.0000	0.0000	0.0008	0.0163	0.1238	0.3467	0.3572	0.1354	0.0189	0.0010	0.0000
5	2.72	0.0000	0.0000	0.0006	0.0136	0.1111	0.3335	0.3681	0.1495	0.0223	0.0012	0.0000

Figure 9

The Vol Surface on March 20, 2011

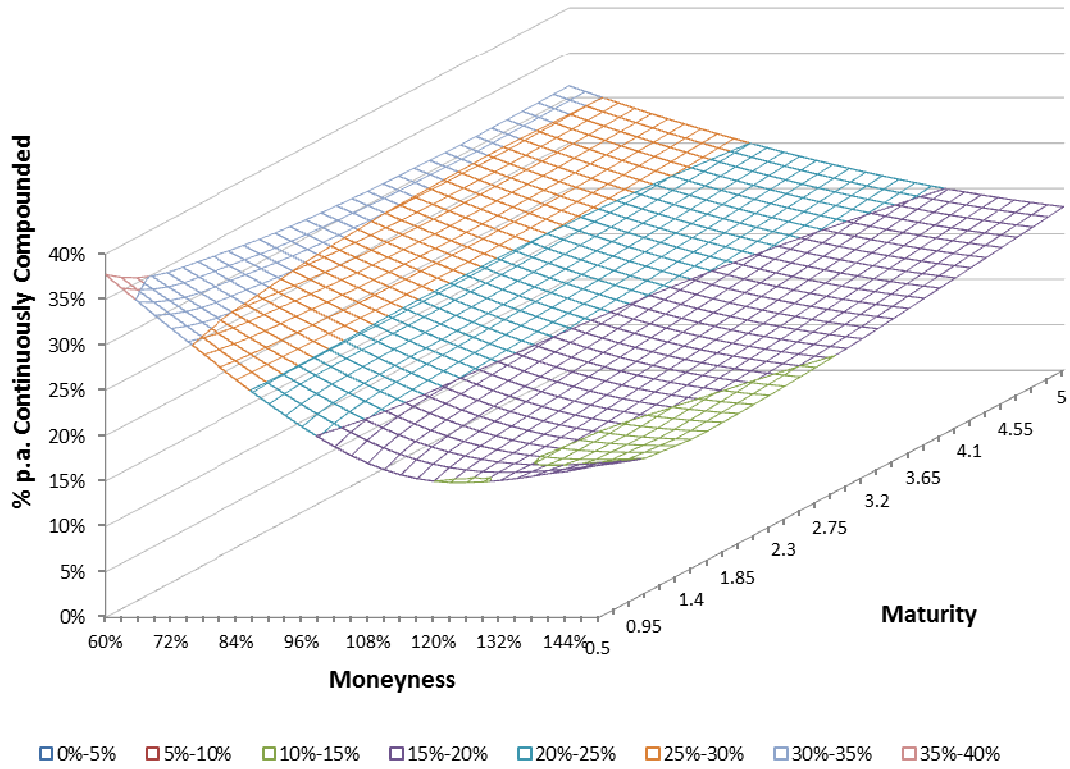


Table 12

**State Prices and the State Space Price Transition Matrix, P,
on April 27, 2011**

Tenor	0.2500	0.5000	0.7500	1.0000	1.2500	1.5000	1.7500	2.0000	2.2500	2.5000	2.7500	3.0000
-35%	0.0054	0.0226	0.0383	0.0496	0.0578	0.0636	0.0681	0.0713	0.0734	0.0747	0.0755	0.0758
-29%	0.0066	0.0185	0.0260	0.0301	0.0325	0.0338	0.0345	0.0346	0.0346	0.0346	0.0343	0.0336
-23%	0.0177	0.0410	0.0461	0.0498	0.0513	0.0515	0.0513	0.0504	0.0496	0.0487	0.0476	0.0462
-16%	0.0446	0.0640	0.0726	0.0733	0.0719	0.0698	0.0679	0.0662	0.0636	0.0606	0.0582	0.0564
-8%	0.1641	0.1562	0.1415	0.1282	0.1176	0.1094	0.1024	0.0962	0.0905	0.0854	0.0807	0.0764
0%	0.4775	0.3023	0.2338	0.1976	0.1732	0.1555	0.1412	0.1295	0.1196	0.1109	0.1032	0.0963
9%	0.2762	0.3157	0.2783	0.2452	0.2192	0.1979	0.1796	0.1643	0.1514	0.1401	0.1297	0.1205
19%	0.0066	0.0701	0.1286	0.1554	0.1661	0.1672	0.1637	0.1582	0.1518	0.1448	0.1374	0.1302
30%	0.0000	0.0024	0.0158	0.0358	0.0553	0.0718	0.0846	0.0937	0.0997	0.1033	0.1051	0.1054
41%	0.0000	0.0001	0.0009	0.0038	0.0092	0.0171	0.0264	0.0358	0.0449	0.0533	0.0610	0.0672
54%	0.0000	0.0000	0.0000	0.0000	0.0002	0.0004	0.0007	0.0011	0.0016	0.0022	0.0028	0.0033
interest factor	0.9988	0.9988	0.9928	0.9819	0.9679	0.9517	0.9327	0.9113	0.8876	0.8623	0.8355	0.8077

The State Price Transition Matrix, P												
	Sigmas	-5	-4	-3	-2	-1	0	1	2	3	4	5
Sigmas	$S_0 \setminus S_T$	0.37	0.45	0.55	0.67	0.82	1.00	1.22	1.49	1.82	2.23	2.72
-5	0.37	0.6711	0.2407	0.0529	0.0048	0.0013	0.0006	0.0009	0.0005	0.0005	0.0004	0.0000
-4	0.45	0.2803	0.3959	0.2447	0.0541	0.0038	0.0001	0.0004	0.0000	0.0000	0.0000	0.0000
-3	0.55	0.0494	0.2244	0.3937	0.2484	0.0560	0.0044	0.0002	0.0000	0.0000	0.0000	0.0000
-2	0.67	0.0057	0.0436	0.2177	0.3903	0.2503	0.0565	0.0033	0.0000	0.0000	0.0000	0.0000
-1	0.82	0.0060	0.0067	0.0410	0.2108	0.3853	0.2490	0.0537	0.0025	0.0000	0.0000	0.0000
0	1.00	0.0054	0.0066	0.0177	0.0446	0.1641	0.4775	0.2762	0.0066	0.0000	0.0000	0.0000
1	1.22	0.0005	0.0005	0.0009	0.0038	0.0399	0.2040	0.3820	0.2508	0.0580	0.0049	0.0001
2	1.49	0.0010	0.0008	0.0012	0.0020	0.0057	0.0421	0.2040	0.3733	0.2431	0.0546	0.0043
3	1.82	0.0015	0.0010	0.0012	0.0017	0.0026	0.0062	0.0408	0.1947	0.3612	0.2323	0.0571
4	2.23	0.0005	0.0003	0.0004	0.0005	0.0007	0.0009	0.0033	0.0352	0.1866	0.3474	0.3130
5	2.72	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0003	0.0321	0.1813	0.8746

Table 13

The Recovered Pricing Kernel and the Natural Probability Transition Matrix

The Natural Probability Transition Matrix, F												
	Sigmas	-5	-4	-3	-2	-1	0	1	2	3	4	5
Sigmas	$S_0 \setminus S_T$	-35%	-29%	-23%	-16%	-8%	0%	9%	19%	30%	41%	54%
-5	-35%	0.6699	0.2530	0.0607	0.0062	0.0019	0.0010	0.0020	0.0014	0.0018	0.0021	0.0000
-4	-29%	0.2657	0.3952	0.2666	0.0662	0.0054	0.0001	0.0008	0.0000	0.0000	0.0000	0.0000
-3	-23%	0.0429	0.2053	0.3930	0.2785	0.0729	0.0072	0.0003	0.0000	0.0000	0.0000	0.0000
-2	-16%	0.0044	0.0355	0.1934	0.3896	0.2899	0.0813	0.0058	0.0000	0.0000	0.0000	0.0000
-1	-8%	0.0040	0.0047	0.0314	0.1814	0.3846	0.3088	0.0805	0.0047	0.0000	0.0000	0.0000
0	0%	0.0029	0.0038	0.0109	0.0308	0.1318	0.4767	0.3330	0.0100	0.0000	0.0000	0.0000
1	9%	0.0002	0.0002	0.0004	0.0022	0.0265	0.1686	0.3813	0.3137	0.0949	0.0114	0.0005
2	19%	0.0004	0.0003	0.0005	0.0009	0.0030	0.0277	0.1625	0.3727	0.3176	0.1016	0.0127
3	30%	0.0004	0.0003	0.0004	0.0006	0.0011	0.0031	0.0248	0.1485	0.3605	0.3302	0.1302
4	41%	0.0001	0.0001	0.0001	0.0001	0.0002	0.0003	0.0014	0.0189	0.1308	0.3467	0.5013
5	54%	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0001	0.0140	0.1128	0.8731
Kernel =		1.86	1.77	1.62	1.44	1.24	1.00	0.83	0.66	0.50	0.35	0.22

Figure 10

The Pricing Kernel and the Closest Fitted Constant Relative Risk Aversion Kernel

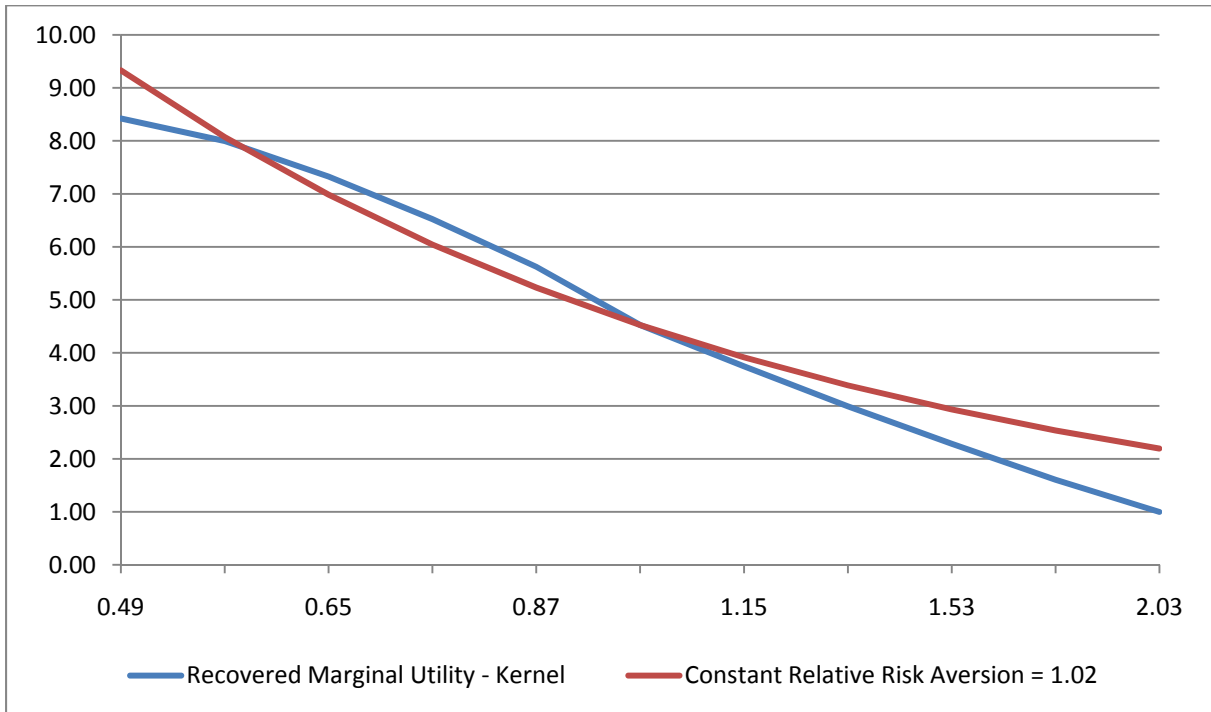


Table 14

The Recovered and the Bootstrapped Natural Marginal Distributions

	Marginal Distributions											
Tenor	0.2500	0.5000	0.7500	1.0000	1.2500	1.5000	1.7500	2.0000	2.2500	2.5000	2.7500	3.0000
-35%	0.0029	0.0120	0.0201	0.0257	0.0295	0.0322	0.0342	0.0358	0.0368	0.0377	0.0385	0.0391
-29%	0.0038	0.0103	0.0143	0.0164	0.0175	0.0180	0.0183	0.0183	0.0183	0.0184	0.0184	0.0183
-23%	0.0109	0.0250	0.0278	0.0296	0.0301	0.0300	0.0296	0.0291	0.0286	0.0283	0.0279	0.0274
-16%	0.0308	0.0438	0.0491	0.0489	0.0474	0.0456	0.0441	0.0428	0.0412	0.0395	0.0383	0.0376
-8%	0.1318	0.1242	0.1110	0.0993	0.0900	0.0829	0.0771	0.0722	0.0681	0.0645	0.0615	0.0590
0%	0.4767	0.2986	0.2277	0.1901	0.1647	0.1465	0.1322	0.1208	0.1117	0.1042	0.0978	0.0925
9%	0.3330	0.3765	0.3274	0.2849	0.2519	0.2252	0.2030	0.1852	0.1709	0.1589	0.1484	0.1398
19%	0.0100	0.1047	0.1897	0.2264	0.2392	0.2384	0.2319	0.2235	0.2146	0.2058	0.1971	0.1892
30%	0.0000	0.0047	0.0306	0.0683	0.1042	0.1341	0.1569	0.1733	0.1845	0.1921	0.1973	0.2005
41%	0.0000	0.0002	0.0024	0.0102	0.0247	0.0455	0.0697	0.0943	0.1184	0.1413	0.1629	0.1822
54%	0.0000	0.0000	0.0000	0.0002	0.0007	0.0016	0.0029	0.0047	0.0068	0.0093	0.0120	0.0145
	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000

	Summary statistics (annualized)											
	0.2500	0.5000	0.7500	1.0000	1.2500	1.5000	1.7500	2.0000	2.2500	2.5000	2.7500	3.0000
mean	0.0463	0.0502	0.0546	0.0566	0.0577	0.0582	0.0581	0.0573	0.0562	0.0549	0.0533	0.0517
sigma	0.1247	0.1441	0.1494	0.1520	0.1534	0.1546	0.1553	0.1551	0.1543	0.1532	0.1518	0.1501
risk free	0.0048	0.0024	0.0096	0.0185	0.0264	0.0336	0.0406	0.0476	0.0544	0.0611	0.0675	0.0738
E - r	0.0415	0.0478	0.0450	0.0381	0.0313	0.0247	0.0175	0.0098	0.0018	-0.0062	-0.0142	-0.0221
Sharpe	0.3330	0.3317	0.3011	0.2508	0.2039	0.1595	0.1124	0.0630	0.0117	-0.0405	-0.0934	-0.1474
ATM vol	0.1453	0.1669	0.1771	0.1820	0.1853	0.1881	0.1909	0.1934	0.1958	0.1982	0.2007	0.2031
	Summary Historical Statistics (Monthly S&P 500 returns from 1960 - 2010 (annualized))											
mean				0.1034								
sigma				0.1547								
risk free				0.0545								
E - r				0.0489								
Sharpe				0.3159								

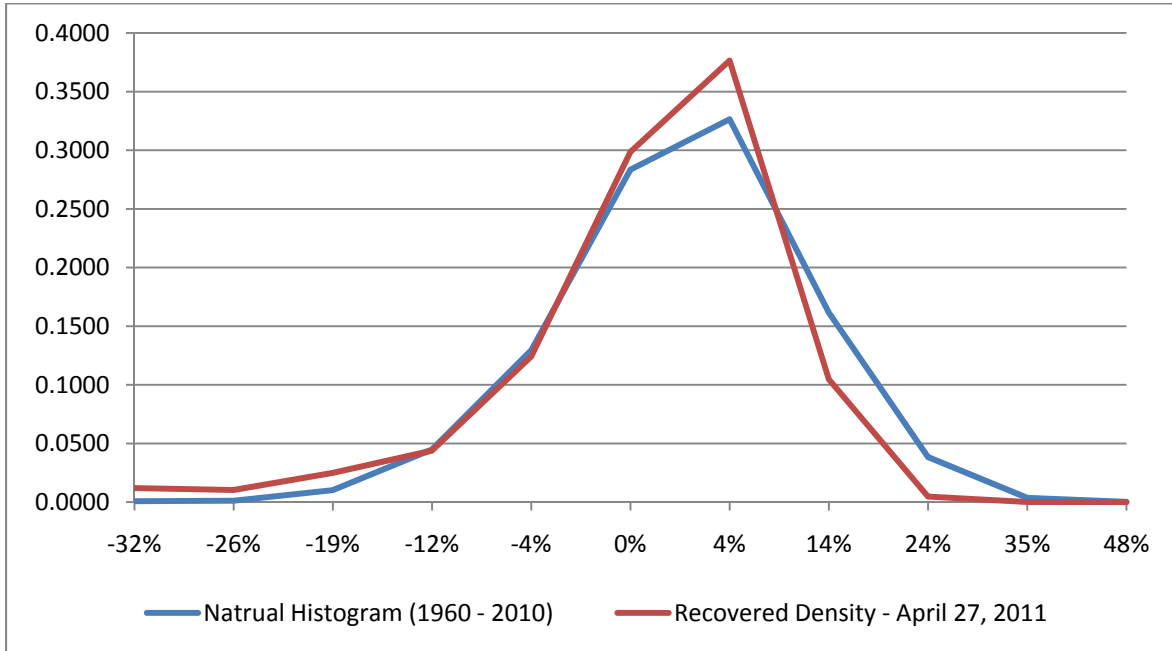
Table 15

The Densities and the Cumulative Distributions for the Recovered and the Bootstrapped Natural Probabilities (six month horizon)

Range	Densities:		Distribution Functions:	
	Bootstrapped	Recovered	Bootstrapped	Recovered
-32%	0.0008	0.0120	0.0008	0.0120
-26%	0.0012	0.0103	0.0020	0.0223
-19%	0.0102	0.0250	0.0122	0.0473
-12%	0.0448	0.0438	0.0570	0.0912
-4%	0.1294	0.1242	0.1864	0.2153
0%	0.2834	0.2986	0.4698	0.5139
4%	0.3264	0.3765	0.7962	0.8904
14%	0.1616	0.1047	0.9578	0.9951
24%	0.0384	0.0047	0.9962	0.9998
35%	0.0036	0.0002	0.9998	1.0000
48%	0.0002	0.0000	1.0000	1.0000

Figure 11

The Recovered and the Bootstrapped Natural Densities



Bibliography

- Barro, R. (2006), Rare Disasters and Asset Markets in the Twentieth Century, *Quarterly Journal of Economics* 121:823-866
- Bernardo, A. and Ledoit, O. (2000), Gain, Loss, and Asset Pricing, *Journal of Political Economy* 108, 144-172
- Breedon, D. and Litzenberger, R. (1978), Prices of State Contingent Claims Implicit in Option Prices, *Journal of Business* 51, 621-651
- Cochrane, J., (2000) Beyond Arbitrage: good Deal Asset Price Bounds in Incomplete Markets, *Journal of Political Economy* 108 (February 2000), 79-119
- Cox, J., and Ross, S. (1976), The Valuation of Options for Alternative Stochastic Processes, *Journal of Financial Economics*, 3, 1976, 145-166
- Cox, J. and Ross, S. (1976) A Survey of Some New Results in Financial Option Pricing Theory, *Journal of Finance* 31, No. 1, May 1976, 383-402
- Cox, J., Ross, S. and Rubinstein, M. (1979), Option Pricing A simplified Approach, *Journal of Financial Economics* 7, 229-263
- Derman, E. and Kani, I. (1994), The Volatility Smile and Its Implied Tree, Goldman Sachs Discussion Paper, January
- Derman, E. and Kani, I. (1998), Stochastic Implied Trees: Arbitrage Pricing with Stochastic Term and Strike Structure of Volatility, *International Journal of Theoretical and Applied Finance* 1: 7-22.
- Dybvig, P., and Rogers, C. (1997), Recovery of Preferences from Observed Wealth in a Single Realization, *Review of Financial Studies* 10, 151-174.
- Dybvig, P. and Ross, S. (1987), Arbitrage, *The New Palgrave Dictionary of Economics*, First Edition, ed by J. Eatwell, M. Milgate, and P. Newman
- Dybvig, P. and Ross, S. (2003), Arbitrage, state prices and portfolio theory, *Handbook of the Economics of Finance*, G.M. Constantinides and M. Harris and R.M. Stulz (ed.), Elsevier, vol 1, chapter 10, 605-637.
- Dupire, B. (1994), Pricing with a Smile, *Risk* 7: 18-20.
- Fama, E. (1970), Efficient Capital Markets: A Review of Theory and Empirical Work, *Journal of Finance*, 25: 383-417.

Hansen, L. and Jagannathan, R. (1991), Implications of Security Market Data for Models of Dynamic Economies, *Journal of Political Economy* 99 (2).

Jackwerth, J. (1999), Optional-Implied Risk-Neutral Distributions and Implied Binomial Trees: A Literature Review, *Journal of Finance* 51: 1611-1631.

Jackwerth, J. and Rubinstein, M. (1996), Recovering Probability Distributions from Option Prices, *Journal of Finance* 51: 1611-1631.

LeRoy, S. and Porter, R. (1981), The Present Value Relation: Tests Based on Implied Variance Bounds, *Econometrica*, 49: 555—574.

Mehra, R. and Prescott, E. (1985), The Equity Premium: A Puzzle, *Journal of Monetary Economics*, 15:145—161.

Merton R.C. (1973), Theory of Rational Option Pricing, *The Bell Journal of Economics and Management Science*, Vol. 4, No. 1: 141-183.

Meyer, Carl (2000), *Matrix Analysis and Applied Linear Algebra*, SIAM, ISBN 0-89871-454-0

Rietz, T. (1988), The Equity Premium: a Solution, *Journal of Monetary Economics* 22: 117-131

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

Ross, S. (2005), *Neoclassical Finance*, Princeton University Press, Princeton, NJ

Ross, S. (1976), The Arbitrage Theory of Capital Asset Pricing, *Journal of Economic Theory* 13, No. 3, December 1976, 341-360

Ross, S. (1976), The Valuation of Options for Alternative Stochastic Processes
Rubinstein, M. (1994), Implied Binomial Trees, *Journal of Finance* 49: 771-818.

Shiller, R. (1981), Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?, *American Economic Review*, 71:421—436.

Weitzman, M. (2007), Subjective Expectations and Asset-Return Puzzles, *American Economic Review*: 97, no. 4, 1102-1130