The Journal of FINANCE
The Journal of THE AMERICAN FINANCE ASSOCIATION

THE JOURNAL OF FINANCE • VOL. LXX, NO. 2 • APRIL 2015

The Recovery Theorem

STEVE ROSS*

ABSTRACT

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

FINANCIAL MARKETS PRICE SECURITIES with payoffs extending out in time, and the hope that they can be used to forecast the future has long fascinated both scholars and practitioners. Nowhere is this more apparent than for the fixed income markets, with an enormous literature devoted to examining the predictive content of forward rates. However, with the exception of foreign exchange and some futures markets, a similar line of research has not developed for other markets. This absence is most notable for the equity markets.

While there exists a rich market in equity options and a well-developed theory of how to use their prices to extract the martingale or risk-neutral probabilities (see Cox and Ross (1976a, 1976b)), there has been a theoretical hurdle to using these probabilities to forecast the probability distribution of future returns, that is, real or natural probabilities. 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 return under the risk-neutral measure is the return under the natural measure with the risk premium subtracted out. The risk premium is a function of both risk and the market's risk aversion, and, thus, to use risk-neutral prices to estimate 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 both the agent's risk aversion and the agent's subjective probability distribution, and neither is directly

*Ross is with the Sloan School, MIT. I want to thank the participants in the UCLA Finance workshop for their insightful comments as well as Richard Roll, Hanno Lustig, Rick Antle, Andrew Jeffrey, Peter Carr, Kevin Atteson, Jessica Wachter, Ian Martin, Leonid Kogan, Torben Andersen, John Cochrane, Dimitris Papanikolaou, William Mullins, Jon Ingersoll, Jerry Hausman, Andy Lo, Steve Leroy, George Skiadopoulos, Xavier Gabaix, Patrick Dennis, Phil Dybvig, Will Mullins, Nicolas Caramp, Rodrigo Adao, Steve Heston, Patrick Dennis, the referee, Associate Editor, and the Editor. All errors are my own. I also wish to thank the participants in the AQR Insight Award and AQR for its support.

DOI: 10.1111/jofi.12092

observable. Instead, we infer them from fitting or "calibrating" market models. Unfortunately, efforts to empirically measure the aversion to risk have led to more controversy than consensus. For example, measures of the coefficient of aggregate risk aversion range from two or three to 500 depending on the model and the macro data used. Additionally, financial 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 (Mehra and Prescott (1985))—which has even given rise to worrisome investment advice based on the view that stocks are uniformly superior to bonds. These conundrums have led some to propose that finance has its equivalent to the dark matter that cosmologists posit to explain their models' behavior for the universe when observables seem insufficient. The dark matter of finance is the very low probability of a catastrophic event and the impact that changes in that perceived probability can have on asset prices (see, for example, Barro (2006) and Weitzmann (2007)). Apparently, however, 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 one month to be higher than the probability calculated from historical 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 has occurred historically or the market requires a very high risk premium to insure against a decline. Without knowing which is the case, it is impossible to separate the two and infer the market's forecast of the event probability.

Determining 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, thus, it has long been argued that any tests of efficient market hypotheses are simultaneously tests of both a particular asset pricing model and the efficient market hypothesis (Fama (1970)). However, if we knew the kernel, we could estimate the variation in the risk premium (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 market hypothesis.

A related issue is the inability to find the current market forecast of the expected return on equities. Unable to obtain this directly from prices as we do with forward rates,[1] we are left to using historical returns and opinion polls of economists and investors, asking them to reveal their estimated risk premiums. It certainly does not 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 does not appear in the resulting formula.

But all is not quite so hopeless. While quite different, the results in this paper are in the spirit of Dybvig and Rogers (1997), who showed that if stock returns follow a recombining tree (or diffusion), then we can reconstruct the agent's

[1] Although these too require a risk adjustment.

utility function from an agent's observed portfolio choice along a single path. Borrowing their nomenclature, we call these results recovery theorems as well. Section I presents the basic analytic framework tying the state price density to the kernel and the natural density. Section II derives the Recovery Theorem, which allows us to estimate the natural probability of asset returns and the market's risk aversion—the kernel—from the state price transition process alone. To do so, two important nonparametric assumptions are introduced in this section. Section III derives the Multinomial Recovery Theorem, which offers an alternative route for recovering the natural distribution for binomial and multinomial processes. Section IV examines the application of these results to some examples and highlights important limitations of the approach. Section V estimates the state price densities at different horizons from S&P 500 option prices on a randomly chosen recent date (April 27, 2011), 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 to the histogram of historical stock returns. In particular, we shed some light on the dark matter of finance by highlighting the difference between the odds of a catastrophe as derived from observed state prices and the odds obtained from historical data. The analysis of Section V is meant to be illustrative and is far from the much needed empirical analysis, but it provides the first use of the Recovery Theorem to estimate the natural density of stock returns. Section VI outlines a model-free test of efficient market hypotheses. Section VII concludes the paper, and points to future research directions.

I. The Basic Framework

Consider a discrete-time world with asset payoffs $g(\theta)$ 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, that is, Arrow-Debreu (Arrow (1952), Debreu (1952)) contingent claims prices, $p(\theta)$ (or in general spaces, a price 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, p_g, of an asset paying $g(\theta)$ in one period is given by

$$p_g = \int g(\theta)\, dP(\theta). \tag{1}$$

Since the sum of the contingent claims prices is the current value of a dollar for sure in the future, letting $r(\theta^0)$ denote the riskless rate as a function of the current state, θ^0, we can rewrite this in the familiar form

$$\begin{aligned} p_g &= \int g(\theta)\, dP(\theta) = (\textstyle\int dP(\theta)) \displaystyle\int g(\theta) \frac{dP(\theta)}{\int dP(\theta)} \\ &\equiv e^{-r(\theta^0)T} \int g(\theta)\, d\pi^*(\theta) \equiv e^{-r(\theta^0)T} E^*[g(\theta)] = E[g(\theta)\phi(\theta)], \end{aligned} \tag{2}$$

where an asterisk denotes the expectation in the martingale measure and where the pricing kernel, that is, the state price/probability, $\phi(\theta)$, is the Radon-Nikodym (see Gurevich and Shilov (1978)) derivative of $P(\theta)$ with respect to the natural measure, which we will denote as $F(\theta)$. With continuous distributions, $\phi(\theta) = p(\theta)/f(\theta)$, where $f(\theta)$ is the natural probability, that is, the actual or relevant subjective probability distribution, and the risk-neutral probabilities, are given by $\pi^*(\theta) = \frac{p(\theta)}{\int p(\theta)d\theta} = e^{r(\theta^0)T}p(\theta)$.

Let θ_i denote the current state and θ_j 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, and thus the stock price can be written as $S(\theta_i)$. From the forward equation for the martingale probabilities we have

$$Q(\theta_i, \theta_j, T) = \int_{\theta} Q(\theta_i, \theta, t)Q(\theta, \theta_j, T - t)d\theta, \tag{3}$$

where $Q(\theta_i, \theta_j, T)$ is the forward martingale probability transition function for going from state θ_i to state θ_j in T periods and where the integration is over the intermediate state θ at time t. Notice that the transition function depends on the time interval and is independent of calendar time.

This is a very general framework that allows for many interpretations. For example, the state could be composed of parameters that describe the motion of the process, for example, the volatility of returns, σ, as well as the current stock price, S, that is, $\theta = (S,\sigma)$. If the distribution of martingale returns is determined only by the volatility, then a transition could be written as a move from $\theta_i = (S,\sigma)$ to $\theta_j = \left(S(1+R), \sigma'\right)$ where R is the rate of return and

$$Q(\theta_i, \theta_j, t) = Q((S, \sigma), (S(1+R), \sigma'), t). \tag{4}$$

To simplify notation we use state prices rather than the martingale probabilities so that we do not have to be continually correcting for the interest factor. Defining the state prices as

$$P(\theta_i, \theta_j, t, T) \equiv e^{-r(\theta_i)(T-t)}Q(\theta_i, \theta_j, T - t) \tag{5}$$

and assuming a time homogeneous process where calendar time is irrelevant, for the transition from any time t to $t+1$, we have

$$P(\theta_i, \theta_j) = e^{-r(\theta_i)}Q(\theta_i, \theta_j). \tag{6}$$

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

$$\phi(\theta_i, \theta_j) = \frac{p(\theta_i, \theta_j)}{f(\theta_i, \theta_j)}, \tag{7}$$

and an equivalent statement of NA is that a positive kernel exists.

A canonical example of this framework is an intertemporal model with a representative agent with additively time-separable preferences and a constant discount factor, δ. We use this example to motivate our results but it is not necessary for the analysis that follows. Letting $c(\theta)$ denote consumption at time t as a function of the state, over any two periods the agent seeks to maximize

$$\max_{\{c(\theta_i),\{c(\theta)\}_{\theta\in\Omega}\}} \{U(c(\theta_i)) + \delta \int U(c(\theta)) f(\theta_i,\theta)\, d\theta\} \tag{8}$$

subject to

$$c(\theta_i) + \int c(\theta) p(\theta_i,\theta) d\theta = w.$$

The first-order condition for the optimum allows us to interpret the kernel as

$$\phi(\theta_i,\theta_j) = \frac{p(\theta_i,\theta_j)}{f(\theta_i,\theta_j)} = \frac{\delta U'(c(\theta_j))}{U'(c(\theta_i))}. \tag{9}$$

Equation (9) 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. The kernel is the agent's marginal rate of substitution as a function of aggregate consumption (see Dybvig and Ross (1987, 2003)).

Notice, too, that in this example the pricing kernel depends only on the marginal rate of substitution between future and current consumption. This path independence is a key element of the analysis in this paper, and the kernel is assumed to have the form of (9), that is, it is a function of the ending state and it depends on the beginning state only through dividing to normalize it.

DEFINITION 1: *A kernel is transition independent if there is a positive function of the states, h, and a positive constant δ such that, for any transition from θ_i to θ_j, the kernel has the form*

$$\phi(\theta_i,\theta_j) = \delta \frac{h(\theta_j)}{h(\theta_i)}. \tag{10}$$

The intertemporally additive utility function is a common example that generates a transition-independent kernel but there are many others.[2]

Using transition independence we can rewrite (7) as

$$p(\theta_i,\theta_j) = \phi(\theta_i,\theta_j) f(\theta_i,\theta_j) = \delta \frac{h(\theta_j)}{h(\theta_i)} f(\theta_i,\theta_j), \tag{11}$$

[2] For example, it is easy to show that Epstein-Zin (1989) recursive preferences also produce a transition independent kernel. Also, see Heston (2004) who uses a similar path-independence assumption to derive the risk-neutral probabilities from the natural probabilities.

where $h(\theta) = U'(c(\theta))$ in the representative agent model. Assuming that we observe the state price transition function, $p(\theta_i, \theta_j)$, our objective is to solve this system to recover the three unknowns: the natural probability transition function, $f(\theta_i, \theta_j)$, the kernel, $\phi(\theta_i, \theta_j) = \delta h(\theta_j)/h(\theta_i)$, and the discount rate, δ. Transition independence, or some variant, is necessary to allow us to separately determine the kernel and the natural probability distribution from equation (7). With no restrictions on the kernel, $\phi(\theta_i, \theta_j)$, or the natural distribution, $f(\theta_i, \theta_j)$, it would not be possible to identify them separately from knowledge of the product alone, $p(\theta_i, \theta_j)$. Roughly speaking, there are more unknowns on the right-hand side of (7) than equations.

An extensive literature provides a variety of approaches to solving this problem. For example, Jackwerth and Rubinstein (1996) and Jackwerth (2000) use implied binomial trees to represent the stochastic process. Ait-Sahalia and Lo (2000) combine state prices derived from option prices with estimates of the natural distribution to determine the kernel. Bliss and Panigirtzoglou (2004) assume constant relative or absolute risk aversion preferences and estimate the elasticity parameter by comparing the predictions of this form with historical data. Bollerslev and Tederov (2011) use high-frequency data to estimate the premium for jump risk in a jump diffusion model and, implicitly, the kernel. These approaches have a common feature: they use the historical distribution of returns to estimate the unknown kernel and thereby link the historical estimate of the natural distribution to the risk-neutral distribution or they make parametric assumptions on the utility function of a representative agent (and often assume that the distribution follows a diffusion).

In the next section, we take a different tack and show that the equilibrium system of equations, (11), can be solved without using either historical data or any assumptions other than a transition independent kernel.

II. The Recovery Theorem

To gain some insight into equation (11) and to position the apparatus for empirical work, from now on we specialize it to a discrete state space model, and, while it is not necessary, we illustrate the analysis with the representative agent formulation

$$U_i' p_{ij} = \delta U_j' f_{ij}, \tag{12}$$

where we can interpret

$$U_i' \equiv U'(c(\theta_i)). \tag{13}$$

But, more generally, U' is any positive function of the state. Writing this in terms of the kernel and denoting the current state θ_i as state $i = 1$,

$$\phi_j \equiv \phi(\theta_1, \theta_j) = \delta(U_j'/U_1'). \tag{14}$$

We define the states from the filtration 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. Notice that, while marginal utility is monotone declining in consumption, it need not be monotone declining in the asset value, $S(\theta_i)$.

Rewriting the state equations (11) in matrix form we have

$$DP = \delta FD, \tag{15}$$

where P is the $m \times m$ matrix of state contingent Arrow-Debreu (1952) prices, p_{ij}, F is the $m \times m$ matrix of the natural probabilities, f_{ij}, and D is the diagonal matrix with the *undiscounted kernel*, that is, the marginal rates of substitution, $\varphi_{\mathrm{j}}\ /\delta$, on the diagonal,

$$D = \left(\frac{1}{U_1'}\right)\begin{bmatrix} U_1' & 0 & 0 \\ 0 & U_i' & 0 \\ 0 & 0 & U_m' \end{bmatrix} = \begin{bmatrix} \phi_1 & 0 & 0 \\ 0 & \phi_i & 0 \\ 0 & 0 & \phi_m \end{bmatrix}\left(\frac{1}{\delta}\right). \tag{16}$$

With a discrete or compact state space for prices, we have to make sure that the model does not permit arbitrage. In a model with exogenous consumption the absence of arbitrage is a simple consequence of an equilibrium with positive state prices, which ensures that the carrying cost net of the dividend compensates for any position that attempts to profit from the increase from the lowest asset value or the decrease from the highest value.

Continuing with the analysis, recall that we observe the state prices, P, and our objective is to see what, if anything, we can infer about the natural measure, F, and the pricing kernel, that is, the marginal rates of substitution. Solving (15) for F as a function of P, we have

$$F = \left(\frac{1}{\delta}\right) DPD^{-1}. \tag{17}$$

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, that is, the risk adjustment, there appears to be no way to close the system and solve for the natural measure, F. Fortunately, however, since F is a matrix whose rows are transition probabilities, it is a stochastic matrix, that is, a positive matrix whose rows sum to one, and there is an additional set of m constraints,

$$Fe = e, \tag{18}$$

where e is a vector with “1” in all the entries.

Using this condition we have

$$Fe = \left(\frac{1}{\delta}\right) DPD^{-1}e = e, \tag{19}$$

or

$$Pz = \delta z, \tag{20}$$

where

$$z \equiv D^{-1}e. \tag{21}$$

This is a characteristic root problem and offers some hope that the solution set is discrete and not an arbitrary cone. With one further condition, the theorem below verifies that this is so and provides us with a powerful result. From NA, P is nonnegative and we will also assume that it is irreducible, that is, all states are attainable from all other states in k steps. For example, if P is positive then it is irreducible. More generally, though, even if there is a zero in the ij entry, it could be possible to get to j in, say, two steps by going from i to k and then from k to j or along a path with k steps. A matrix P is irreducible if there is always some path such that any state j can be reached from any state i.[3]

THEOREM 1 (The Recovery Theorem): *If there is NA, if the pricing matrix is irreducible, and if it is generated by a transition independent kernel, then there exists a unique (positive) solution to the problem of finding the natural probability transition matrix,* F, *the discount rate,* δ, *and the pricing kernel,* ϕ. *In other words, for any given set of state prices there is a unique compatible natural measure and a unique pricing kernel.*

Proof: Existence can also be proven directly, but it follows immediately from the fact that P is assumed to be generated from F and D as shown above. The problem of solving for F is equivalent to finding the characteristic roots (eigenvalues) and characteristic vectors (eigenvectors) of P since, if we know δ and z such that

$$Pz = \delta z, \tag{22}$$

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

From the Perron Frobenius Theorem (see Meyer (2000)) all nonnegative irreducible matrices have a unique positive characteristic vector, z, and an associated positive characteristic root, λ. The characteristic root $\lambda = \delta$ is the subjective rate of time discount. Letting z denote the unique positive characteristic vector with root λ, we can solve for the kernel as

$$\frac{U'(c(\theta_i))}{U'(c(\theta_1))} = \left(\frac{1}{\delta}\right)\phi_i = d_{ii} = \frac{1}{z_i}. \tag{23}$$

To obtain the natural probability distribution, from our previous analysis,

$$F = \left(\frac{1}{\delta}\right) DPD^{-1} \tag{24}$$

[3] Notice that, since the martingale measure is absolutely continuous with respect to the natural measure, P is irreducible if F is irreducible.

and

$$f_{ij} = \left(\frac{1}{\delta}\right)\frac{\phi_i}{\phi_j}p_{ij} = \left(\frac{1}{\delta}\right)\frac{U_i^{'}}{U_j^{'}}p_{ij} = \left(\frac{1}{\lambda}\right)\frac{z_j}{z_i}p_{ij}. \tag{25}$$

Q.E.D.

Notice that if the kernel is not transition independent then we have no assurance that the probability transition matrix can be separated from the kernel as in the proof. Notice, too, that there is no assurance that the kernel will be monotone in the ordering of the states by, for example, stock market values.[4]

COROLLARY 1: *The subjective discount rate, δ, is bounded above by the largest interest factor.*

Proof: From The Recovery Theorem the subjective rate of discount, δ, is the maximum characteristic 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, the row sums of P are the interest factors and the maximum row sum is the maximum interest factor. Q.E.D.

Now let's turn to the case in which 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, that is, pricing is risk-neutral.*

Proof: In this case we have

$$Pe = \gamma e, \tag{26}$$

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

$$F = \left(\frac{1}{\gamma}\right)P. \tag{27}$$

Q.E.D.

Given the apparent ease of creating intertemporal models satisfying the usual assumptions without risk-neutrality, this result may seem strange, but it is a consequence of having a finite irreducible process for state transition. When we extend the recovery result to multinomial processes that are unbounded, this is no longer the case.

[4] In an earlier draft it was shown that the Recovery Theorem also holds if there is an absorbing state.

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

THEOREM 3: *The risk-neutral 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

$$\frac{\pi^*\left(\theta_i, \theta_j\right)}{f\left(\theta_i, \theta_j\right)} = \frac{e^{r(\theta_i)} p\left(\theta_i, \theta_j\right)}{f\left(\theta_i, \theta_j\right)} = e^{r(\theta_i)} \phi\left(\theta_i, \theta_j\right) = e^{r(\theta_i)} \frac{\delta U'\left(c\left(\theta_j\right)\right)}{U'\left(c\left(\theta_i\right)\right)}, \tag{28}$$

we know that the ratio is declining in $c\left(\theta_j\right)$. Fixing θ_i, since both densities integrate to one, defining v by $e^{r(\theta_i)}\delta U'(v) = U'(c(\theta_i))$, it follows that $\pi^* > f$ for $c < v$ and $\pi^* < f$ for $c > v$. 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. Q.E.D.

COROLLARY 2: *In a one-period world the market displays a risk premium, that is, 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, the return in the risk-neutral measure is

$$R^* \sim R - Z + \varepsilon, \tag{29}$$

where R is the natural return, Z is strictly nonnegative and ε is mean zero conditional on $R - Z$. Taking expectations we have

$$E[R] = r + E[Z] > r. \tag{30}$$

Q.E.D.

The Recovery Theorem embodies the central intuitions of recovery and is sufficiently powerful for the subsequent empirical analysis. However, before leaving this section we should note that, while there are extensions to continuous state spaces, the Recovery Theorem as developed here relies heavily on the finiteness of the state space. In the next section, we take a different tack and derive a recovery theorem when the state space is infinite and generated by a binomial or multinomial process, and in Section IV we examine a continuous state space example.

III. A Binomial and Multinomial Recovery Theorem

While the Recovery Theorem can be applied to a binomial or multinomial process, doing so requires a truncation of the state space. To avoid this step and since such processes are so ubiquitous in finance (see Cox, Ross, and Rubinstein (1979)), 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 are wholly consumed. Tree growth is governed by a multinomial process and the state of the economy is $<\theta_i, H>$, $i = 1, \ldots, m$. 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, H, and not on the complete state, $<\theta_i, H>$. The value of the tree—the market value of the economy's assets—is given by $S = S(\theta_i, H)$. Since tree height and therefore consumption follow a multinomial process, S also follows a multinomial, but in general jump sizes change with the state.

The marginal utility of consumption depends only on the dividend, and without loss of generality we set initial $U'(kH) = 1$. The equilibrium equations are

$$p_{ij}(H) = \delta U'(ka_j H) f_{ij}, \tag{31}$$

or, in terms of the *undiscounted kernel* $\varphi_j = U'_j$

$$p_{ij}(H) = \delta \phi_j f_{ij}. \tag{32}$$

In matrix notation,

$$P = \delta FD, \tag{33}$$

$$F = \left(\frac{1}{\delta}\right) PD^{-1}, \tag{34}$$

and, since F is a stochastic matrix,

$$Fe = \left(\frac{1}{\delta}\right) PD^{-1}e = e, \tag{35}$$

or

$$PD^{-1}e = \delta e. \tag{36}$$

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

$$\left(\frac{1}{\delta}\right) D^{-1}e = P^{-1}e, \tag{37}$$

and F is recovered as

$$F = \left(\frac{1}{\delta}\right) PD^{-1}. \tag{38}$$

We can now proceed node by node and recover F and δD, but the analysis does not recover δ and ϕ separately. However, taking advantage of the recombining feature of the process we can recover δ and ϕ separately. For simplicity, consider a binomial process that jumps to a or b. The binomial is recurrent, that is, it eventually returns arbitrarily close to any starting position, which is equivalent to irreducibility in this setting. For a binomial, the infinite matrix has only two nonzero elements in any row, and at a particular node we 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 S and there is no exact path that returns to S, then we can get arbitrarily close to S along a path where the number of up (a) steps, i, and the number of down (b) steps, $n-i$, satisfy

$$\frac{i}{n-i} \to -\frac{\log b}{\log a} \tag{39}$$

for large n.

Sparing the obvious continuity analysis, we simply assume that the binomial recurs in two steps, that is, $ab = 1$. At the return step from aH to H, then, since the current state is $<\theta_a, aH>$, the price of receiving one in one period is

$$p_{ab}(aH) = \delta\left(\frac{U'(kH)}{U'(kaH)}\right) f_{ab} = \delta\left(\frac{1}{\phi_a}\right) f_{ab}. \tag{40}$$

Since we have recovered $\delta\phi_a$ from equation (37) we can now solve separately for δ and ϕ_a. The analysis is similar for the general multinomial case.

To implement recovery, if the current state is a, say, we need to know $p_{ba}(H)$ and $p_{bb}(H)$, and if there are no contingent forward markets that allow them to be observed directly, we can compute them from current prices. The prices of going from the current state to a or b in three steps along the paths *(a,b,a)* and *(a,b,b)* when divided by the price of returning to the current state in two steps by the path *(a,b)* are $p_{ba}(H)$ and $p_{bb}(H)$, respectively. Alternatively, if we know the current price of returning to the current state in two steps, $p_{a\cdot 1}$, then

$$\begin{gathered} p_{a\cdot 1} = p_{aa}(H)\, p_{ab}(aH) + p_{ab}(H)\, p_{ba}(bH) \\ = \delta^2\left[\phi_a f_{aa}\left(\tfrac{1}{\phi_a}\right)(1 - f_{aa}) + \phi_b(1 - f_{aa})\left(\tfrac{1}{\phi_b}\right)(1 - f_{bb})\right] \\ = \delta^2(1 - f_{aa})(1 + f_{aa} - f_{bb}). \end{gathered} \tag{41}$$

is an independent equation that completes the system and allows it to be solved for δ, F, and ϕ.

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

$$Pe = \gamma e. \tag{42}$$

Hence,

$$\left(\frac{1}{\delta}\right) D^{-1} e = P^{-1} e = \left(\frac{1}{\gamma}\right) e, \tag{43}$$

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

If P is not of full rank, while there is a solution to

$$Fe = \left(\frac{1}{\delta}\right) PD^{-1}e = e, \tag{44}$$

in general there is a (nonlinear) subspace of potential solutions with dimension equal to the rank of P, and we cannot uniquely recover the kernel and the probability matrix. As an example, consider a simple binomial process that jumps to a with probability f. In this case P has two identical rows and recombining gives us a total of three equations in the four unknowns δ, f, ϕ_a and ϕ_b:

$$p_a = \delta \phi_a f, \tag{45}$$

$$p_b = \delta \phi_b (1 - f), \tag{46}$$

and

$$p_{a.1} = 2\delta^2 f (1 - f), \tag{47}$$

which, with positivity, has a one-dimensional set of solutions.

However, in the special case where the interest rate is state independent even if the matrix is of less than full rank, 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 and 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 transition matrix is of less than full rank, then we can restrict the potential solutions, but recovery is not unique. If the state prices are independent of the state, then risk-neutrality is always one possible solution.*

Proof: See the analysis above preceding the statement of the theorem.

Q.E.D.

In Section V, we use the Recovery Theorem but we could also have used the Multinomial Recovery Theorem. Which approach is preferable depends on the availability of contingent state prices and, ultimately, is an empirical question. Now we look at some special cases.

A. Relative Risk Aversion

An alternative approach to recovery is to assume a functional form for the kernel. Suppose, for example, that the kernel is generated by a constant relative risk aversion (CRRA) 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

$$p_{xy}(H) = \phi(kH, kyH) f_{xy}. \tag{48}$$

Hence, after the current dividend, the value of stock (the tree) is

$$S(a,H) = p_{aa}(H)\left[S(a,aH) + kaH\right] + p_{ab}(H)\left[S(b,bH) + kbH\right] \tag{49}$$

and

$$S(b,H) = p_{ba}(H)\left[S(a,aH) + kaH\right] + p_{bb}(H)\left[S(b,bH) + kbH\right]. \tag{50}$$

Assuming CRRA,

$$\phi(x,y) = \delta\left(\frac{y}{x}\right)^{-\gamma}, \tag{51}$$

this system is linear with the solution

$$S(x,H) = \gamma_x H, \tag{52}$$

where

$$\begin{pmatrix}\gamma_a\\ \gamma_b\end{pmatrix} = \begin{bmatrix} 1-\delta f_a a^{1-\gamma} & -\delta(1-f_a)b^{1-\gamma} \\ -\delta(1-f_b)a^{1-\gamma} & 1-\delta f_b b^{1-\gamma}\end{bmatrix}^{-1}\begin{pmatrix}\delta f_a k a^{1-\gamma} + \delta(1-f_a)kb^{1-\gamma}\\ \delta(1-f_b)ka^{1-\gamma} + \delta f_b k b^{1-\gamma}\end{pmatrix}. \tag{53}$$

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,

$$S(a,H) = \gamma_a H \rightarrow \gamma_a aH = S(a,aH),\, or\ \gamma_b bH = S(b,bH) \tag{54}$$

and

$$S(b,H) = \gamma_b H \rightarrow \gamma_a aH = S(a,aH),\, or\ \gamma_b bH = S(b,bH). \tag{55}$$

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

Without making use of recombination, the state price equations for this system are given by

$$p_{aa} = \delta f_a a^{-\gamma}, \tag{56}$$

$$p_{ab} = \delta(1-f_a)b^{-\gamma}, \tag{57}$$

$$p_{ba} = \delta(1-f_b)a^{-\gamma}, \tag{58}$$

and

$$p_{bb} = \delta f_b b^{-\gamma}. \tag{59}$$

Assuming state independence, $f_b \neq 1 - f_a$, these are four independent equations in the four unknowns δ, γ, f_a, f_b, and the solution is given by

$$\begin{pmatrix} f_a \\ f_b \end{pmatrix} = \left(\frac{p_{ab} p_{ba}}{p_{aa} p_{bb}} - 1 \right)^{-1} \begin{pmatrix} \frac{p_{ab}}{p_{bb}} - 1 \\ \frac{p_{ba}}{p_{aa}} - 1 \end{pmatrix}, \tag{60}$$

$$\gamma = \frac{-ln\left(\frac{f_a}{1-f_a}\right) + ln\left(\frac{p_{aa}}{p_{ab}}\right)}{ln\left(\frac{b}{a}\right)}, \tag{61}$$

and

$$\delta = \frac{p_{aa} a^{\gamma}}{f_a}. \tag{62}$$

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

$$p_a(H) = \delta f a^{-\gamma} \tag{63}$$

and

$$p_a(H) = \delta(1 - f) b^{-\gamma}. \tag{64}$$

This cannot be augmented by recombining since, assuming $ab = 1$,

$$p_a(bH) = \delta f \left(\frac{1}{b^{-\gamma}} \right) = \delta f b^{\gamma} = \delta f a^{-\gamma}, \tag{65}$$

which is identical to the first equation. In other words, while the parametric assumption has reduced identifying the two-element kernel to recovering a single parameter, γ, it has also eliminated one of the equations. As we have shown, however, 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

$$p_{aa} = \delta f_a a^{-\gamma}, \tag{66}$$

$$p_{ab} = \delta(1 - f_a) \left(\frac{k_a}{k_b} \right)^{\gamma} b^{-\gamma}, \tag{67}$$

$$p_{ba} = \delta(1 - f_b) \left(\frac{k_b}{k_a} \right)^{\gamma} a^{-\gamma}, \tag{68}$$

and

$$p_{bb} = \delta f_b b^{-\gamma}. \tag{69}$$

These are four independent equations that can be solved for the four unknowns $\delta, \gamma, f_a, and\, f_b$.

IV. An Example, Comments, and Extensions

Consider a model with a lognormally distributed payoff at time T and a representative agent with a CRRA utility function,

$$U(S_T) = \frac{S_T^{1-\gamma}}{1-\gamma}. \tag{70}$$

The future stock payoff, the consumed dividend, is lognormal,

$$S_T = e^{\left(\mu - \frac{1}{2}\sigma^2\right)T + \sigma\sqrt{T}z}, \tag{71}$$

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

The pricing kernel is given by

$$\phi_T = \frac{e^{-\delta T} U'(S_T)}{U'(S)} = e^{-\delta T}\left[\frac{S_T}{S}\right]^{-\gamma}, \tag{72}$$

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

$$P_T(S, S_T) = \phi_T\left(\frac{S_T}{S}\right) n_T(lnS_T) = e^{-\delta T}\left[\frac{S_T}{S}\right]^{-\gamma} n\left(\frac{lnS_T - \left(\mu - \frac{1}{2}\sigma^2\right)T}{\sigma\sqrt{T}}\right)\left(\frac{1}{S_T}\right), \tag{73}$$

where $n(\cdot)$ is the normal density function, or, in terms of the log of consumption, $s \equiv ln(S)$ and $s_T \equiv ln(S_T)$,

$$P_T(s, s_T) = e^{-\delta T} e^{-\gamma(s_T - s)} n\left(\frac{s_T - \left(\mu - \frac{1}{2}\sigma^2\right)T}{\sigma\sqrt{T}}\right). \tag{74}$$

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 I displays natural transition probability matrix, F, the pricing kernel, and the matrix P of transition prices. The units of relative stock movement, S_T/S, are units of sigma on a grid from -5 to $+5$. While sigma can be chosen as the standard deviation of the derived martingale measure from P, we chose the current at-the-money implied volatility from option prices on the S&P 500 index as of March 15, 2011.

Table I

Fixed Lognormally Distributed Future Payoff and a Constant Relative Risk Aversion($\gamma = 3$) Pricing Kernel

The matrices below are derived from the one-period model presented in Section IV. The rows and columns in the matrices refer to ranges for the stock price state variable, for example, three standard deviations from the current level is 1.82. The Sigma = 0 row is the current state.

Panel A: The State Space Transition Matrix (P)

	Sigmas	−5	−4	−3	−2	−1	0	1	2	3	4	5
Sigmas	$S_0 \backslash S_T$	0.37	0.45	0.55	0.67	0.82	1	1.22	1.49	1.82	2.23	2.72
−5	0.37	0.000	0.000	0.001	0.005	0.015	0.019	0.008	0.001	0.000	0.000	0.000
−4	0.45	0.000	0.000	0.001	0.008	0.028	0.034	0.015	0.003	0.000	0.000	0.000
−3	0.55	0.000	0.000	0.002	0.015	0.051	0.062	0.028	0.005	0.000	0.000	0.000
−2	0.67	0.000	0.000	0.003	0.028	0.092	0.113	0.051	0.008	0.001	0.000	0.000
−1	0.82	0.000	0.000	0.006	0.051	0.168	0.205	0.092	0.015	0.001	0.000	0.000
0	1	0.000	0.000	0.010	0.092	0.306	0.374	0.168	0.028	0.002	0.000	0.000
1	1.22	0.000	0.001	0.019	0.168	0.558	0.681	0.306	0.051	0.003	0.000	0.000
2	1.49	0.000	0.001	0.034	0.306	1.016	1.241	0.558	0.092	0.006	0.000	0.000
3	1.82	0.000	0.003	0.062	0.558	1.852	2.262	1.016	0.168	0.010	0.000	0.000
4	2.23	0.000	0.005	0.113	1.016	3.374	4.121	1.852	0.306	0.019	0.000	0.000
5	2.72	0.000	0.008	0.205	1.852	6.148	7.509	3.374	0.558	0.034	0.001	0.000
Kernel, $\phi =$	20.086	11.023	6.05	3.32	1.822	1	0.549	0.301	0.165	0.091	0.05	

Panel B: The Natural Probability Transition Matrix (F)

	Sigmas	−5	−4	−3	−2	−1	0	1	2	3	4	5
Sigmas	$S_0 \backslash S_T$	0.37	0.45	0.55	0.67	0.82	1	1.22	1.49	1.82	2.23	2.72
−5	0.37	0.000	0.000	0.002	0.028	0.171	0.381	0.312	0.094	0.010	0.000	0.000
−4	0.45	0.000	0.000	0.002	0.028	0.171	0.381	0.312	0.094	0.010	0.000	0.000
−3	0.55	0.000	0.000	0.002	0.028	0.171	0.381	0.312	0.094	0.010	0.000	0.000
−2	0.67	0.000	0.000	0.002	0.028	0.171	0.381	0.312	0.094	0.010	0.000	0.000
−1	0.82	0.000	0.000	0.002	0.028	0.171	0.381	0.312	0.094	0.010	0.000	0.000
0	1	0.000	0.000	0.002	0.028	0.171	0.381	0.312	0.094	0.010	0.000	0.000
1	1.22	0.000	0.000	0.002	0.028	0.171	0.381	0.312	0.094	0.010	0.000	0.000
2	1.49	0.000	0.000	0.002	0.028	0.171	0.381	0.312	0.094	0.010	0.000	0.000
3	1.82	0.000	0.000	0.002	0.028	0.171	0.381	0.312	0.094	0.010	0.000	0.000
4	2.23	0.000	0.000	0.002	0.028	0.171	0.381	0.312	0.094	0.010	0.000	0.000
5	2.72	0.000	0.000	0.002	0.028	0.171	0.381	0.312	0.094	0.010	0.000	0.000

With an assumed market return of 8%, and a standard deviation of 20% we calculate the characteristic vector of P. As anticipated, there is one positive vector that exactly equals the pricing kernel shown in Table I, and the characteristic root is $e^{-0.02} = 0.9802$, as was assumed. Solving for the natural transition matrix, F, we have exactly recovered the posited lognormal density.

This static example fits the assumptions of the Recovery Theorem closely except for having a continuous rather than a discrete distribution. The closeness of the results with the actual distribution and kernel suggests that applying the theorem by truncating the tail outcomes is an appropriate approach in this case. 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.

Finding this result in a continuous space example is important since the Recovery Theorem was proven on a discrete and, therefore, bounded state space. To explore the impact of significantly loosening this assumption, we can extend the example to allow for consumption growth.

Assuming that consumption follows a lognormal growth process,

$$S_T = S_0 e^{(\mu - \frac{1}{2}\sigma^2)T + \sigma\sqrt{T}z}, \tag{75}$$

state prices are given by

$$P_T(s, s_T) = e^{-\delta T} e^{-\gamma(s_T - s)} n\left(\frac{s_T - s - \left(\mu - \frac{1}{2}\sigma^2\right)T}{\sigma\sqrt{T}}\right). \tag{76}$$

Taking logs,

$$\log\, P_T(x, y) = -\delta T - \gamma(s_T - s) - \left(\frac{1}{2\sigma^2 T}\right)\left(s_T - s - \left(\mu - \frac{1}{2}\sigma^2\right)T\right)^2 \\ -\log\sqrt{2\pi T}\sigma, \tag{77}$$

and as $(s_T - s)$ varies, state prices depend on the quadratic form

$$-\left(\frac{1}{2\sigma^2 T}\right)(s_T - s)^2 - \left(\gamma - \left(\frac{1}{\sigma^2}\right)\left(\mu - \frac{1}{2}\sigma^2\right)\right)(s_T - s) \\ -\left(\delta T + \left(\frac{T}{2\sigma^2}\right)\left(\mu - \frac{1}{2}\sigma^2\right)^2 + \log\sqrt{2\pi T}\sigma\right). \tag{78}$$

Since the prices follow a diffusion, even if we assume that we know σ it is not possible to extract the three parameters μ, γ, and δ from the two relevant parameters of the quadratic,

$$\gamma - \left(\frac{1}{\sigma^2}\right)\left(\mu - \frac{1}{2}\sigma^2\right) \; \mathit{and} \; \left(\delta + \left(\frac{1}{2\sigma^2}\right)\left(\mu - \frac{1}{2}\sigma^2\right)^2\right)T. \tag{79}$$

This indeterminacy first arose with the Black-Scholes (1973) and Merton (1973) option pricing formula and similar diffusion equations for derivative pricing in which, with risk-neutral pricing, the risk-free interest rate is substituted for the drift, μ, in the valuation formulas.

What happens, then, if we attempt a continuous space analogue to the Recovery Theorem? The analogous space characteristic equation to be solved is

$$\int_0^\infty p(s, y)\, v(y)\, dy = \lambda v(s). \tag{80}$$

By construction $v(x) = 1/U'(x)$ and $e^{-\delta T}$ satisfy this equation, but they are not the unique solutions, and a little mathematics verifies that any exponential, $e^{\alpha x}$, also satisfies the characteristic equation with characteristic value

$$\lambda(\alpha) = e^{-\delta T} e^{(\alpha-\gamma)(\mu-\frac{1}{2}\sigma^2)T+\frac{1}{2}\sigma^2 T(\alpha-\gamma)^2}. \tag{81}$$

Since α is arbitrary, this agrees with the earlier finding and the well-established intuition that, given risk-neutral prices and even assuming that σ is observable (as it would be for a diffusion), we cannot determine the mean return, μ, of the underlying process.

Why, then, did we have success in finding a solution in the original static version?[5] One important difference between the two models arises when we discretize by truncation. By truncating the process we are implicitly making the marginal utility the same in all states beyond a threshold, which is a substitute for bounding the process and the state space. A natural conjecture would be that if the generating kernel has a finite upper bound on marginal utility (and, perhaps, a nonzero lower bound as well), then the recovered solution will be unique.[6] Whether the kernel is generated by a representative agent with bounded marginal utility cannot be resolved by theory alone, but a practical approach would be to examine the stability of the solution with different extreme truncations.

A more directly relevant comparison between the two models is that in the growth model the current state has no impact on the growth rate. When combined with a CRRA kernel, the result is that state prices depend only on the difference between the future state and the current state. This makes the growth model a close relative of the state independent binomial process examined in the previous section. As we show, an alternative approach to aid recovery is to introduce explicit state dependence. For example, we could model the dependence of the distribution on a volatility process by taking advantage of the observed strong empirical inverse relation between changes in volatility and current returns. This could once again allow us to apply the Recovery Theorem as above.[7]

V. Applying the Recovery Theorem

With the rich market for derivatives on the S&P 500 index and on futures on the index, we assume that the market is effectively complete along dimensions related to the index, that is, both value and the states of the return process. The Recovery Theorem relies on knowledge of the martingale transition matrix and, given the widespread interest in using the martingale measure for pricing derivative securities, it is not surprising that an extensive literature

[5] From (74) it is easy to see that $e^{\gamma x}$ is the unique characteristic solution.

[6] The multiplicity of solutions in the continuous case was pointed out to me by Xavier Gabaix. Carr and Yu (2012) have established recovery with a bounded diffusion and Ross (2013) has done so with a bounded kernel.

[7] An explicit example is available from the author upon request.

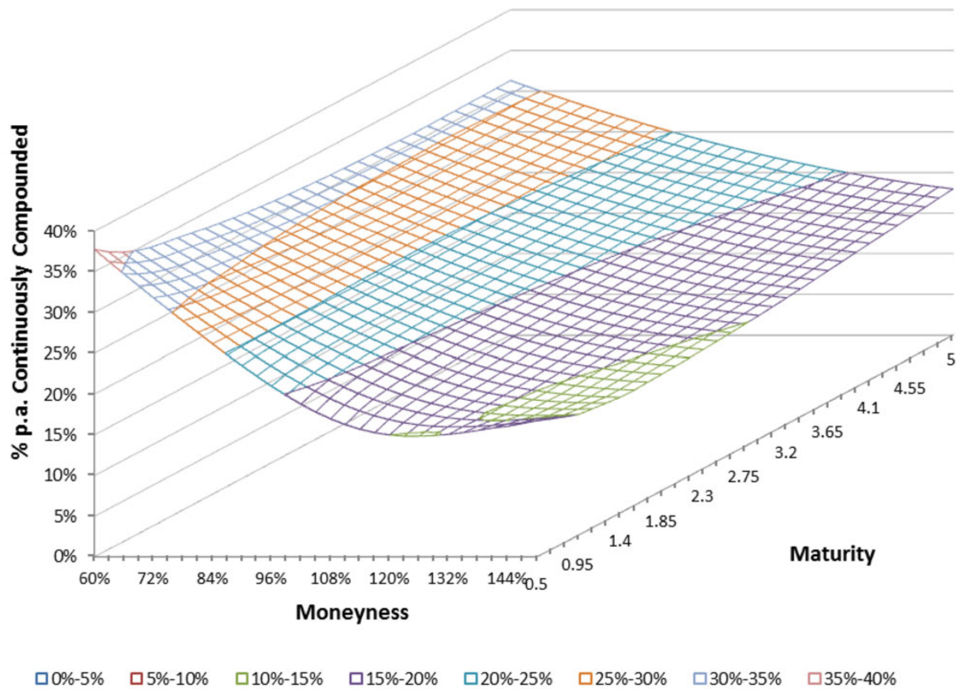

Figure 1. The implied volatility surface on March 20, 2011. The surface of implied volatilities on puts and calls on the S&P500 index on March 20, 2011 is drawn as a function of both time to maturity in years ("tenor") and the strike price divided by current price ("moneyness"). Option prices are typically quoted in terms of implied volatilities from the Black-Scholes (1973) and Merton (1973) formula, and are displayed here on the vertical axis. The source of the data used in this paper is a bank over-the-counter bid/offer sheet.

estimates the martingale measure (see, for example, Rubinstein (1994), Jackwerth and Rubinstein (1996), Jackwerth (2000), Derman and Kani (1994 and 1998), Dupire (1994), Ait-Sahalia and Lo (1998), Figlewski (2008)). We draw on only the most basic findings of this work.

Figure 1 displays the surface of implied volatilities on S&P puts and calls, the "volatility surface," on March 20, 2011, 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 (1973) and Merton (1973) formula, that is, the volatilities that when put into the model give the market premium for the option. Note that doing so is not a statement of the validity of the Black-Scholes (1973) and Merton (1973) model; rather, it is simply a transformation of the market determined premiums into a convenient way to quote them. The source of the data used in this paper is a bank over-the-counter bid/offer sheet. While the data are in broad agreement with exchange traded options, we choose this source since the volume on the over-the-counter market is multiples of that on the exchange even for at the money contracts.[8]

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 problem 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. An explanation for this pattern 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.

Implied volatilities are a function of the risk-neutral probabilities, the product of the natural probabilities and the pricing kernel (that is, risk aversion and time discounting), and as such they embody all of the information needed to determine state prices. Since all contracts can be formed as portfolios of options (Ross (1976a)), it is well known that from the volatility surface and the formula for the value of a call option we can derive the state price distribution, $p(S,T)$ at any tenor T:

$$C(K,T) = \int_0^\infty [S-K]^+ \, p(S,T)\, dS = \int_K^\infty [S-K]\, p(S,T)\, dS, \tag{82}$$

where $C(K,T)$ is the current price of a call option with a strike of K and a tenor of T. Differentiating twice with respect to the strike, we obtain the Breeden and Litzenberger (1978), result that

$$p(K,T) = C''(K,T). \tag{83}$$

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, with state defined by the return from holding the index until T. Setting the grid size of index movements at 0.5%, the S&P 500 call options on April 27, 2011 produced the state prices reported in the top table of Table II. The results are broadly sensible with the exception of the relatively high implied interest rates at longer maturities, which we address below.

To apply the Recovery Theorem, though, we need the $m \times m$ state price transition matrix,

[8] *Bank for International Settlements Quarterly Review*, June 2012 Statistical Annex, pages A135 and A136. While there is some lack of clarity as to the exact option terms, the notional on listed equity index options is given as $197.6 billion of notional, and that for over-the-counter equity options is given as $4.244 trillion.

Table II
State Prices and Recovered Probabilities

Panel A displays the Arrow-Debreu (1952) state prices for the current values of $1 in the relevant stock price return range given in the left-hand column at the tenors given in the top row. These are derived by taking the numerical second derivative with respect to the strikes of traded call option prices from a bank offer sheet. The row labeled “discount factor” sums each column of the first state price matrix to obtain the implied risk-free discount factors. Panel B is the estimated table of contingent state prices that are consistent with the given Arrow-Debreu (1952) state prices. These are derived by applying the forward equation to find the transition matrix that best fits the Arrow-Debreu (1952) state prices subject to the constraint that the resulting transition matrix has unimodal rows. The two top rows and two leftmost columns express the state variable in terms of both standard deviations from the current level and the stock price.

Panel A: State Prices on April 27, 2011

Return\Tenor	0.25	0.5	0.75	1	1.25	1.5	1.75	2	2.25	2.5	2.75	3
−36%	0.005	0.023	0.038	0.050	0.058	0.064	0.068	0.071	0.073	0.075	0.076	0.076
−30%	0.007	0.019	0.026	0.030	0.032	0.034	0.034	0.035	0.035	0.035	0.034	0.034
−24%	0.018	0.041	0.046	0.050	0.051	0.052	0.051	0.050	0.050	0.049	0.048	0.046
−16%	0.045	0.064	0.073	0.073	0.072	0.070	0.068	0.066	0.064	0.061	0.058	0.056
−9%	0.164	0.156	0.142	0.128	0.118	0.109	0.102	0.096	0.091	0.085	0.081	0.076
0%	0.478	0.302	0.234	0.198	0.173	0.155	0.141	0.129	0.120	0.111	0.103	0.096
9%	0.276	0.316	0.278	0.245	0.219	0.198	0.180	0.164	0.151	0.140	0.130	0.120
20%	0.007	0.070	0.129	0.155	0.166	0.167	0.164	0.158	0.152	0.145	0.137	0.130
31%	0.000	0.002	0.016	0.036	0.055	0.072	0.085	0.094	0.100	0.103	0.105	0.105
43%	0.000	0.000	0.001	0.004	0.009	0.017	0.026	0.036	0.045	0.053	0.061	0.067
57%	0.000	0.000	0.000	0.000	0.000	0.000	0.001	0.001	0.002	0.002	0.003	0.003
Discount Factor	1.000	0.993	0.983	0.969	0.953	0.938	0.920	0.900	0.883	0.859	0.836	0.809

Panel B: The State Price Transition Matrix (P)

	Sigmas	−5	−4	−3	−2	−1	0	1	2	3	4	5
Sigmas	$S_0 \backslash S_T$	0.64	0.70	0.76	0.84	0.91	1	1.09	1.20	1.31	1.43	1.57
−5	0.64	0.671	0.241	0.053	0.005	0.001	0.001	0.001	0.001	0.001	0.000	0.000
−4	0.70	0.280	0.396	0.245	0.054	0.004	0.000	0.000	0.000	0.000	0.000	0.000
−3	0.76	0.049	0.224	0.394	0.248	0.056	0.004	0.000	0.000	0.000	0.000	0.000

(Continued)

Table II—*Continued*

Panel B: The State Price Transition Matrix (P)

	Sigmas	−5	−4	−3	−2	−1	0	1	2	3	4	5
−2	0.84	0.006	0.044	0.218	0.390	0.250	0.057	0.003	0.000	0.000	0.000	0.000
−1	0.91	0.006	0.007	0.041	0.211	0.385	0.249	0.054	0.002	0.000	0.000	0.000
0	1.00	0.005	0.007	0.018	0.045	0.164	0.478	0.276	0.007	0.000	0.000	0.000
1	1.09	0.001	0.001	0.001	0.004	0.040	0.204	0.382	0.251	0.058	0.005	0.000
2	1.20	0.001	0.001	0.001	0.002	0.006	0.042	0.204	0.373	0.243	0.055	0.004
3	1.31	0.002	0.001	0.001	0.002	0.003	0.006	0.041	0.195	0.361	0.232	0.057
4	1.43	0.001	0.000	0.000	0.001	0.001	0.001	0.003	0.035	0.187	0.347	0.313
5	1.57	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.032	0.181	0.875

$$P = [p\,(i, j)]\text{, where } p\ (i, j) \text{ is the state } i \text{ price of an Arrow-Debreu security paying off in state } j. \tag{84}$$

Unfortunately, since a rich forward market for options does not exist, and we do not directly observe P, we have to estimate it from the state price distributions at different tenors.

Currently, the system is in some particular state, c, and we observe the current prices of options across strikes and tenors. As shown above in equation (83), from these option prices we can extract the state prices at each future date T,

$$p^T\,(c) = \langle p\,(1, T)\,,\ \ldots, p\,(m, T)\rangle\,. \tag{85}$$

Let the stock price at time T, S_T, index the states and denote the current stock price, S_0. The row of the state price transition matrix, P, corresponding to the current state, c, is simply p_c, that is, the vector of one period ahead state prices with $T = 1$ in equation (85). Since our intention is illustrative, we have ignored the potential state dependence on past returns and on other variables such as implied volatility itself, and identified 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, for example, a quarter, is a good surrogate for the price path—this is clearly a matter for further study.

To solve for the remaining elements of P, we apply the forward equation recursively to create the sequence of $m - 1$ equations:

$$p^{t+1} = p^t P, t = 1,\ \ldots, m - 1, \tag{86}$$

where m is the number of states. Each of the equations in (86) indicates that the current state price for a security paying off in state j at time $T + 1$ is the state price for a payment at time T in some intermediate state k multiplied by the transition price of going from state k to state j, $p(k,j)$, added up over all the possible intermediate states, k. Thus, by looking at only m time periods we have the m^2 equations necessary to solve for the m^2 unknown $p(i,j)$ transition prices.

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. In an effort to minimize the errors in the estimation of P, it was required that the resulting state prices, the rows of P, should be unimodal.

The grid is chosen to be from -5 to $+5$ standard deviations with a standard deviation of 9% per quarter. This seems to be a reasonable compromise between fineness and coverage in the tails. We then implement the analysis above 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 II. The resulting forward transition price matrix, P, is shown in Table II under the table of the state prices, p^t.

The state prices in Table II should sum to the riskless interest factor. The rates are relatively accurate out to about one year but then rise from 1.85% at one year to 7.93% at three years. The three-year rate is significantly higher than three-year (swap) rates at the time and indicative of a bias in the computation of the state prices, which impacts some subsequent results as we point out below. This result has nothing to do with the recovery theory per se, but rather is a consequence of the crudeness in the computation of state prices from option prices and speaks to the critical need to do a better job at this step.

The final step applies the Recovery Theorem to the transition pricing matrix, P, to recover the pricing kernel and the resulting natural probability (quarterly) transition matrix shown in Table III. The kernel declines monotonically as the stock value rises, but this need not be the case. The recovered characteristic root, δ, the social discount rate in a representative agent model, is 1.0018. Alternatively, if we were to use monthly data instead of quarterly observations, the characteristic root is 0.9977, which is less than one—as it should be. This serves as a warning about the sensitive nature of the estimation procedure.

Table IV shows the recovered natural marginal distributions at the future dates, summary statistics for the recovered distributions, and comparable summary statistics for the historical distribution estimated by a bootstrap of S&P 500 returns from 60 years of data (1960 to 2010). Table IV also displays the implied volatilities from the option prices on April 27, 2011. The summary statistics display significant differences between the recovered and historical distributions. For the recovered distribution, which is a forward-looking measure, the annual expected return at all horizons is approximately 6% per year as compared with 10% per year for the historical measure. The recovered standard deviation, on the other hand, is comparable at about 15% per year—an unsurprising 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 and realized returns. The upward biased estimates of the risk-free interest rate beyond two years are the source of the risk premium (and thus the Sharpe ratio) in Table IV declining and turning negative at 2.5 years.

Notice that the at-the-money implied volatilities are significantly higher than those derived from the recovered distribution. This is a phenomenon 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. This difference is consistent with the existence of a risk premium for bearing volatility risk, but in and of itself it is not dispositive.

Table V compares the recovered natural density and distribution with those obtained from a bootstrap of historical data, and Figure 2 plots these densities. Of particular interest is what the results say about the long-standing concern with tail events. Rietz (1988) argues that a large but unobserved probability of a catastrophe—"tail risk"—could explain the equity risk premium puzzle, that is, the apparent dominance of stocks over bonds and related questions. Barro (2006) lends 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) provides a deep theoretical argument in

Table III

The Recovered Pricing Kernel and the Natural Probability Transition Matrix

Applying the Recovery Theorem to the data in Table II, the matrix below displays the resulting natural transition probabilities from the ranges for the stock price returns in the far left column to the identical ranges in the top rows. The bottom row displays the recovered kernel for the given stock ranges in the top row.

The Natural Probability Transition Matrix (F)												
	Sigmas	−5	−4	−3	−2	−1	0	1	2	3	4	5
Sigmas	$S_0\backslash S_T$	−36%	−30%	−24%	−16%	−9%	0%	9%	20%	31%	43%	57%
−5	−36%	0.670	0.253	0.061	0.006	0.002	0.001	0.002	0.001	0.002	0.002	0.000
−4	−30%	0.266	0.395	0.267	0.066	0.005	0.000	0.001	0.000	0.000	0.000	0.000
−3	−24%	0.043	0.205	0.393	0.278	0.073	0.007	0.000	0.000	0.000	0.000	0.000
−2	−16%	0.004	0.035	0.193	0.390	0.290	0.081	0.006	0.000	0.000	0.000	0.000
−1	−9%	0.004	0.005	0.031	0.181	0.385	0.309	0.080	0.005	0.000	0.000	0.000
0	0%	0.003	0.004	0.011	0.031	0.132	0.477	0.333	0.010	0.000	0.000	0.000
1	9%	0.000	0.000	0.000	0.002	0.027	0.169	0.381	0.314	0.095	0.011	0.000
2	20%	0.000	0.000	0.000	0.001	0.003	0.028	0.163	0.373	0.318	0.102	0.013
3	31%	0.000	0.000	0.000	0.001	0.001	0.003	0.025	0.148	0.361	0.330	0.130
4	43%	0.000	0.000	0.000	0.000	0.000	0.000	0.001	0.019	0.131	0.347	0.501
5	57%	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.014	0.113	0.873
Kernel, ϕ =		1.86	1.77	1.62	1.44	1.24	1	0.83	0.66	0.50	0.35	0.22

Table IV

The Recovered and the Bootstrapped Natural Marginal Distributions

Each column of Panel A represents a time horizon (years) and the entries in that column are the recovered probabilities of the respective future S&P 500 ranges. Panel B computes the associated summary statistics for the recovered marginal distribution of each time horizon. Panel C displays the comparable summary statistics derived from monthly S&P 500 returns over the 50-year period from 1960 to 2010.

Panel A: The Recovered Marginal Distributions

Return\Tenor	0.25	0.5	0.75	1	1.25	1.5	1.75	2	2.25	2.5	2.75	3
−36%	0.003	0.012	0.020	0.026	0.030	0.032	0.034	0.036	0.037	0.038	0.038	0.039
−30%	0.004	0.010	0.014	0.016	0.017	0.018	0.018	0.018	0.018	0.018	0.018	0.018
−24%	0.011	0.025	0.028	0.030	0.030	0.030	0.030	0.029	0.029	0.028	0.028	0.027
−16%	0.031	0.044	0.049	0.049	0.047	0.046	0.044	0.043	0.041	0.039	0.038	0.038
−9%	0.132	0.124	0.111	0.099	0.090	0.083	0.077	0.072	0.068	0.065	0.062	0.059
0%	0.477	0.299	0.228	0.190	0.165	0.146	0.132	0.121	0.112	0.104	0.098	0.092
9%	0.333	0.377	0.327	0.285	0.252	0.225	0.203	0.185	0.171	0.159	0.148	0.140
20%	0.010	0.105	0.190	0.226	0.239	0.238	0.232	0.224	0.215	0.206	0.197	0.189
31%	0.000	0.005	0.031	0.068	0.104	0.134	0.157	0.173	0.184	0.192	0.197	0.200
43%	0.000	0.000	0.002	0.010	0.025	0.045	0.070	0.094	0.118	0.141	0.163	0.182
57%	0.000	0.000	0.000	0.000	0.001	0.002	0.003	0.005	0.007	0.009	0.012	0.015

Panel B: Recovered Summary Statistics (Annualized)

	0.25	0.5	0.75	1	1.25	1.5	1.75	2	2.25	2.5	2.75	3
Mean	0.051	0.055	0.06	0.062	0.063	0.063	0.063	0.062	0.061	0.06	0.058	0.057
Sigma	0.117	0.14	0.147	0.15	0.152	0.153	0.154	0.154	0.153	0.152	0.151	0.149
Risk-free	0.005	0.002	0.01	0.018	0.026	0.034	0.041	0.048	0.054	0.061	0.068	0.074
$E-r$	0.047	0.053	0.05	0.043	0.036	0.03	0.022	0.015	0.007	−0.001	−0.009	−0.017
Sharpe	0.399	0.376	0.34	0.287	0.239	0.193	0.146	0.096	0.044	−0.008	−0.061	−0.115
ATM vol.	0.145	0.167	0.177	0.182	0.185	0.188	0.191	0.193	0.196	0.198	0.201	0.203

(Continued)

Table IV—*Continued*

Panel C: Annualized Monthly S&P 500 Returns	
Mean	0.103
Sigma	0.155
Risk-free	0.055
$E - r$	0.049
Sharpe	0.316

Table V

The Densities and the Cumulative Distributions for the Recovered and the Bootstrapped Natural Probabilities

(Six month horizon, bootstrap using data from 1/1/1960 to 11/30/2010)

The rows of the table correspond to ranges for the S&P 500 index for six months from the date April 27, 2011. The first and third columns are from the historical distribution obtained by bootstrapping independent monthly return observations from the period 1960 to 2010. The second and fourth columns display the comparable distribution results from the recovered distribution of Table IV.

	Densities		Distribution Functions	
Range	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

support of fat tails. More pithily, Merton Miller observed after the 1987 crash that 10 standard deviation events seemed to be happening every few years.

As was suggested in the introduction, tail risk is economists' version of 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. By separating the kernel from the forward-looking probabilities embedded in option prices, we can shed some light on the dark matter and estimate the market's probability of a catastrophe. As Figure 2 shows, the recovered density has a fatter left tail than the historical distribution. Table V puts the probability of a six-month decline in excess of 32% at 0.0008, or four in 5,000 bootstraps. By contrast, the recovered density puts this probability at 1.2%. Similarly, the historical probability of a decline in excess of 26% in a six-month period is 0.002 (10 times in 5,000 bootstraps) while the recovered market probability of 0.0223 is 10 times greater, at over 2%.

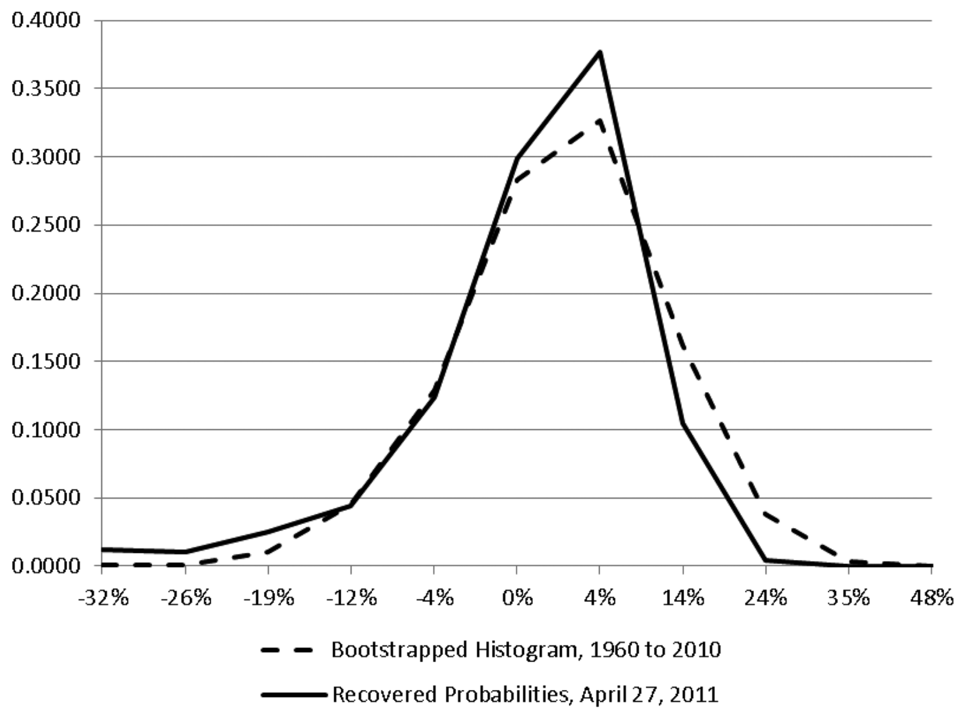

Figure 2. The recovered and the bootstrapped natural densities.

This is only a 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, starting with more carefully estimating the state price density from option prices and then estimating the state price transition matrix from the state price density at different horizons and strikes. Many improvements are also required to accurately recover the kernel and the natural measure implicit in the state prices. Such improvements could certainly alter the implications drawn from the single example analyzed in this section.

VI. 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 (see Fama (1970)). Under the hypothesized conditions of the Recovery Theorem we can separate efficiency from a pricing model and to that extent we can derive model-free tests of the efficient market hypothesis. Ross (2005) proposes an approach to testing efficient market hypotheses that depends on finding an upper bound to the volatility of the pricing kernel; such a bound is a simple byproduct of recovery.

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

$$\sigma(\phi) \geq \left(e^{-rT}\right)\frac{\mu}{\sigma}, \tag{87}$$

where μ is the absolute value of the mean excess return and σ is the standard deviation on any asset, which implies that $\sigma(\phi)$ is bounded from below by the largest observed discounted 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 V we can compute the variance of the kernel, for example, one year out. The computation is straightforward and the resulting variance is

$$\sigma^2(\phi) = 0.1065, \tag{88}$$

or an annual standard deviation of

$$\sigma(\phi) = 0.3264, \tag{89}$$

which, ignoring the small interest factor, is the upper limit for the Sharpe ratio for any strategy to be consistent with efficient markets. It is also a bound used in the literature on when a deal is "too good" (see Cochrane (2000) and Bernardo and Ledoit (2000) for a discussion of good deals, and Ross (1976a) for an early use of the bound for asset pricing).

Alternatively (see Ross (2005)), we can decompose excess returns, x_t, on an asset or portfolio strategy according to

$$x_t = \mu(I_t) + \varepsilon_t, \tag{90}$$

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

$$\sigma^2(x_t) = \sigma^2(\mu(I_t)) + \sigma^2(\varepsilon_t) \leq E\left[\mu^2(I_t)\right] + \sigma^2(\varepsilon_t). \tag{91}$$

Rearranging yields an upper bound on the R^2 of the regression,

$$R^2 = \frac{\sigma^2(\mu(I_t))}{\sigma^2(x_t)} \leq \frac{E\left[\mu^2(I_t)\right]}{\sigma^2(x_t)} \leq e^{2rT}\sigma^2(\varphi), \tag{92}$$

that is, the R^2 is bounded above by the volatility of the pricing kernel (see Ross (2005)). Notice that 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. A potential advantage of tests such as these is that they depend on the second moments, much like the volatility tests of efficiency, and as such might be more robust than standard t-statistic tests on coefficient.

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 with 90% idiosyncratic in an efficient market. Hence, any test of an investment strategy that uses publicly available data and has the ability to predict future returns with $R^2 > 10\%$ would be a violation of efficient markets independent of the specific asset pricing model being used, subject to the maintained assumptions of the Recovery Theorem. Of course, any such strategy must also overcome transactions costs to be an implementable violation—a strategy that could not overcome those costs would be purely of academic interest.

VII. Summary and Conclusions

NA implies the existence of positive Arrow-Debreu (1952) state prices, a risk-neutral measure under which the expected return on any asset is the risk-free rate, and, equivalently, the existence of a strictly positive pricing kernel that can be used to price all assets by taking the expectation of their payoffs weighted by the kernel. To this framework we add some additional nonparametric conditions. First, we make the common assumption that the underlying process is Markov in the state variables, and for implementation we discretized the state space. Second, we assume that the kernel is transition independent, that is, it is a function of the final state and depends only on the current state as a normalization, as is the case for the marginal rate of substitution over time for an agent with an intertemporally additively separable utility function.

In this setting, we are able to prove the Recovery Theorem, which allows us to uniquely determine the kernel, the discount rate, future values of the kernel, and the underlying natural probability distribution of returns from the transition state prices alone. There is no need to use either the historical distribution of returns or independent parametric assumptions on preferences to find the market's subjective distribution of future returns. Put another way, we have a setting in which, even though risk-neutral probabilities are the product of an unknown kernel (that is, risk aversion) and natural probabilities, the two can be disentangled from each other.

A novel element of the approach is that it focuses on the state transition matrix whose elements give the price of one dollar in a future state, conditional on any other state. This is a challenge for implementation when we do not observe the price of a dollar in a future state conditional on being in a different state from the current one, due to the absence of appropriate contingent forward markets. An example illustrates how to find these transition prices from the state prices for different maturities derived from the market prices of simple options by using a version of the forward equation for Markov processes. The accuracy with which this can be done and the accuracy with which state prices can be estimated from option prices will eventually determine how useful the Recovery Theorem will be both empirically and practically. In an example we assume that the state could be summarized by the current level of the index.

This is clearly not the case: for example, implied volatility is also a relevant state variable. Extending the empirical analysis to include such variables will be important, along with gauging the extent to which this has a significant impact. Particularly for short horizons, this remains to be explored.

Finding the limitations and appropriate extensions of the Recovery Theorem is a rich research agenda. Several conjectured extensions to allow recovery include bounding the assumed kernel, bounding the underlying process, and incorporating various forms of state dependence in the process. In general, we want to know what is necessary to apply the theorem or extensions to continuous or unbounded processes, and what sort of bounds on the underlying process or bounds on the assumed kernel will allow recovery. We also need to further explore the Multinomial Recovery Theorem and perhaps introduce weak parametric assumptions into both recovery theorems. While we have focused on the equity markets, bounds on the process are natural for interest rates and fixed income markets, and this will be an important area to explore (see Carr and Yu (2012)).

Once we have recovered the kernel (that is, the market's risk aversion) and the market's subjective assessment of the distribution of returns, these can be used in a host of applications. We can use the market's future distribution of returns much as we use forward rates as forecasts of future spot rates, albeit without a theoretical bias. Institutional asset holders, such as pension funds, use historical estimates of the risk premium on the market as an input in asset allocation models. The market's current subjective forecast should be superior, and at the least is of interest. Project valuation also uses historical estimates of the risk premium. Risk control models, such as VAR (Value at Risk), typically use historical estimates to determine the risk of various books of business and this too would be enhanced by using the recovered distribution. Moreover, with time-series data we will be able to test these predictions against realizations.

The above results can also be applied to a wide variety of markets, such as fixed income, currency, and futures. Indeed, beyond using forward rates, we make little use of interest rate options to estimate the future probability distribution of rates, and applying recovery techniques to this market is a promising line of research. For the stock market, the kernel and the recovered distribution can be used to recover the distribution of returns for individual stocks and to examine the host of market anomalies and potential violations of market efficiency. The ability to better assess the market's perspective of the likelihood of a catastrophic decline 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, and 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 would also shed light on the controversy as to whether the market is too volatile to be consistent with rational pricing models (see, for example, Leroy and Porter (1981), Shiller (1981)).

In conclusion, contrary to finance folklore, under the appropriate assumptions it is possible to separate risk aversion from the natural distribution, and estimate each from market prices. With a pun intended, we have only scratched the surface of discovering the forecasts imbedded in market prices both for the market itself and, more generally, for the economy as a whole.

Initial submission: November 8, 2011; Final version received: June 28, 2013
Editor: Campbell Harvey

REFERENCES

Ait-Sahalia, Yacine, and Andrew W. Lo, 1998, Nonparametric estimation of state price densities implicit in financial asset prices, *Journal of Finance* 53, 499–547.

Ait-Sahalia, Yacine, and Andrew W. Lo, 2000, Nonparametric risk management and implied risk aversion, *Journal of Econometrics* 94, 9–51.

Arrow, Kenneth, 1952, An extension of the basic theorems of classical welfare economics, in J. Neyman, ed.: *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability* (University of California Press, Berkeley).

Barro, Robert J., 2006, Rare disasters and asset markets in the twentieth century, *Quarterly Journal of Economics* 121, 823–866.

Bernardo, Antonio E., and Olivier Ledoit, 2000, Gain, loss, and asset pricing, *Journal of Political Economy* 108, 144–172.

Black, Fischer, and Myron Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–654.

Bliss, Robert. R, and Nikolaos Panigirtzoglou, 2004, Option-implied risk aversion estimates, *Journal of Finance* 59, 407–446.

Bollerslev, Tim, and Viktor Todorov, 2011, Tails, fears, and risk premia, *Journal of Finance* 66, 2165–2211.

Breeden, Douglas T., and Robert Litzenberger, 1978, Prices of state contingent claims implicit in option prices, *Journal of Business* 51, 621–651.

Carr, Peter, and Jiming Yu, 2012, Risk, return, and Ross recovery, *Journal of Derivatives* 20, 38–59.

Cochrane, John H., 2000, Beyond arbitrage: Good deal asset price bounds in incomplete markets, *Journal of Political Economy* 108, 79–119.

Cox, John C., and Stephen A. Ross, 1976a, The valuation of options for alternative stochastic processes, *Journal of Financial Economics* 3, 145–166.

Cox, John C., and Stephen A. Ross, 1976b, A survey of some new results in financial option pricing theory, *Journal of Finance* 31, 383–402.

Cox, John C., Stephen A. Ross, and Mark E. Rubinstein, 1979, Option pricing: A simplified approach, *Journal of Financial Economics* 7, 229–263.

Debreu, Gerard, 1952, *Theory of Value* (John Wiley & Sons, New York).

Derman, Emanuel, and Iraj Kani, 1994, The volatility smile and its implied tree, *Goldman Sachs Research Notes*.

Derman, Emanuel, and Iraj Kani, 1998, Stochastic implied trees: Arbitrage pricing with stochastic term and strike structure of volatility, *International Journal of Theoretical and Applied Finance* 1, 7–22.

Dupire, Bruno, 1994, Pricing with a smile, *Risk* 7, 18–20.

Dybvig, Philip H., and Christopher G. Rogers, 1997, Recovery of preferences from observed wealth in a single realization, *Review of Financial Studies* 10, 151–174.

Dybvig, Philip H., and Stephen A. Ross, 1987, Arbitrage, in John Eatwell, Murray Milgate, and Peter Newman, eds.: *The New Palgrave Dictionary of Economics, First Edition* (Macmillan, Stockton, CA).

Dybvig, Philip H., and Stephen A. Ross, 2003, Arbitrage, state prices and portfolio theory, in George M. Constantinedes, Milton Harris, and Rene M. Stultz, eds.: *Handbook of the Economics of Finance* (Elsevier, New York, Amsterdam, and London).

Epstein, Larry G., and Stanley E. Zin, 1989, Substitution, risk aversion, and the temporal behavior of consumption growth and asset returns: A theoretical framework, *Econometrica* 57, 937–969.

Fama, Eugene F., 1970, Efficient capital markets: A review of theory and empirical work, *Journal of Finance* 25, 383–417.

Figlewski, Stephen, 2008, Estimating the implied risk neutrality density for the U.S. market portfolio, in Tim Bollerslev, Jeffrey R. Russell, and Mark Watson, eds.: *Volatility and Time Series Economics: Essays in Honor of Robert F. Engle,* Chapter 8 (Oxford University Press, Oxford, UK).

Gurevich, B. L., and G.E. Shilov, 1978, Richard A. Silverman, trans.: *Integral, Measure, and Derivative: A Unified Approach* (Dover Publications, New York).

Hansen, Lars P., and Ravi Jagannathan, 1991, Implications of security market data for models of dynamic economies, *Journal of Political Economy* 99, 557–590.

Heston, Steve, 2004, Option valuation with infinitely divisible distributions, *Quantitative Finance* 4, 515–524.

Jackwerth, Jens C., 2000, Recovering risk aversion from option prices and realized returns, *Review of Financial Studies* 13, 4333–4451.

Jackwerth, Jens C., and Mark Rubinstein, 1996, Recovering probability distributions from option prices, *Journal of Finance* 51, 1611–1631.

LeRoy, Stephen F., and Richard D. Porter, 1981, The present value relation: Tests based on implied variance bounds, *Econometrica* 49, 555–574.

Mehra, Rajnish, and Edward C. Prescott, 1985, The equity premium: A puzzle, *Journal of Monetary Economics* 15, 145–161.

Merton, Robert C., 1973, Theory of rational option pricing, *Bell Journal of Economics and Management Science* 4, 141–183.

Meyer, Carl D., 2000, *Matrix Analysis and Applied Linear Algebra*, (SIAM, Philadelphia, PA).

Rietz, Thomas A., 1988, The equity premium: A solution, *Journal of Monetary Economics* 22, 117–131.

Ross, Stephen A., 1976a, Options and efficiency, *Quarterly Journal of Economics* 90, 75–89.

Ross, Stephen A., 1976b, The arbitrage theory of capital asset pricing, *Journal of Economic Theory* 13, 341–360.

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

Ross, Stephen A., 2013, The Recovery Theorem in a continuous space model, Working paper.

Rubinstein, Mark E., 1994, Implied binomial trees, *Journal of Finance* 49, 771–818.

Shiller, Robert J., 1981, Do stock prices move too much to be justified by subsequent changes in dividends?, *American Economic Review* 71, 421–436.

Weitzman, Martin L., 2007, Subjective expectations and asset-return puzzles, *American Economic Review* 97, 1102–1130.