Ross Recovery Summary

By Edward Mehrez

Overview

State prices are the product of risk aversion—the pricing kernel—and the natural probability distribution. From derivatives prices we can observe distribution of state prices. The question then becomes: can we separate the markets probability distribution over returns and risk aversion? For example, we would like to know, given forward rates, how much of the rate is due to the market's forecast and how much is a risk premium. In models with a representative agent, this is equivalent to knowing both the agent's risk aversion and the agent's subjective probability distribution, neither of which is observable but instead, inferred from calibrating market models.

Ross finds a decomposition of state prices into risk aversion and the natural probability distribution assuming that the pricing kernel is irreducible and transition independent (defined below). He calls this a recovery theorem. He proves his recovery theorem in two settings: finite state space, and a multinomial (potentially countably infinite) state space. The decomposition then allows us to express the market transition probabilities in terms of the stochastic discount factor and state prices, both of which can be estimated. The proof of the finite state space is essentially a fairly straight forward application of the Perron-Frobenius theorem. The proof of the result in the multinomial setting requires an additional assumption and follows from an unclear proof by (implicit) induction.

As a corollary to the recovery theorem, the subjective discount rate is shown to be bounded above by the largest interest factor. In addition, in the finite state case, if the riskless rate is state independent, then pricing is shown to be risk neutral. This is a **very surprising** result and in fact does not hold true for the countably infinite multinomial case. At the moment, I do not have good intuition for this result, but Ross claims that it is an artifact of having a finite irreducible process for a state transition (see bottom of pg 623).

Ross then goes on to demonstrate the recovery theory in two different ways. First, he shows for a "static" example, that given both the utility function (CRRA in his example) and the stock price distribution (lognormal) that using the recovery theorem, one can recover the natural probability measure using the pricing kernel (SDF) and the state prices. In the given example, the utility function and the price distribution are used to derive the expressions for the SDF and the state prices (through the Black-Scholes-Merton formula). He then calibrates the standard deviation and mean return of the price process (price distribution) to market data and shows that the resulting recovered transition probability distribution coincides with the lognormal distribution. This section (Section IV) proves to be deeper than just a verification of the theorem for an example. First, Ross points out that although the theorem was proven for the discrete and multinomial cases, it still seems to recover with a continuous distribution when considering a static problem (moving from one known initial state to an unknown state). However, once the dynamic problem is considered where one first transitions from a known state to an unknown state, then from the unknown state to another unknown state (from time 0 information set perspective), problems arise in that no implicit truncation of the distribution can be used. Details on this will be given below.

Ross Recovery Summary

By Edward Mehrez

In section V, the recovery theorem is applied to the S&P 500 index to recover the market transition probabilities on April 27, 2011. What is remarkable about the recovery method is that it doesn't need a training set. That is no historical data is explicitly needed to recover the distribution. The only place that Ross makes use of historical data is for the sake of comparison. Using historical returns, he constructs a bootstrapped histogram of returns and compares it to the recovered histogram of returns.

In the last section, a "model-free" test of the efficient market hypothesis (EMH) is proposed that essentially bounds the R^2 one can get from a factor model that would still be consistent with the EMH. Thus, any test of an investment strategy that uses publically available information 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, such a strategy must also overcome transactions costs to really be a violation. His bound doesn't take into account any transactions costs.

Basic Framework (§ II-III)

The basic framework is a discrete-time world with asset payoffs $g(\theta)$ at time T, contingent on the state $\in \Omega$. From the fundamental theorem of asset pricing (FTAC), no arbitrage implies the existence of a positive state space prices, $p(\theta)$ (or in more general spaces, a price distribution function $P(\theta)$). The current value p_g of an asset paying $g(\theta)$ in period T is given by

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

Where an asterisk denotes the expectation with respect to the martingale measure and where the pricing kernel, $\phi(\theta)$ is the Radon-NIkodym derivative of $P(\theta)$ wrt the natural (subjective) measure, $F(\theta)$ or in the case of continuous distributions, $\phi(\theta) = p(\theta)/f(\theta)$, where f is the subjective pdf. The risk neutral probabilities are given by $\pi^*(\theta) = \frac{p(\theta)}{\int p(\theta) d\theta} = e^{r(\theta_0)T} p(\theta)$.

Our **first assumption** is that the **asset value follows a Markov process**. Ignoring the effects of the time value of money temporarily, the (martingale) transition density function, Q, must satisfy the Chapman-Kolmogorov equation:

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

for any $t < T$. The idea behind (3) is that the probability of transitioning from state θ_i to θ_j from time 0 to time T can be decomposed into the total probability of first transitioning to some arbitrary state θ at the intermediate time t and then transitioning to the state θ_j at time T. Furthermore, since Q is a martingale measure, we also get time homogeneity which is why we can view everything from the time zero perspective, i.e. instead of thinking about the transition density from t to T we can just think about the transition density from 0 to T-t.

Taking back into account the time value of money, and making the **second assumption** that the process is **time homogenous**, the state price distribution is given by

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

Under the Markov assumption and assuming a continuous distribution for illustrative purposes, the kernel (state price per unit of probability (density)) as

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

In the familiar world of a representative agent with additive time-separable preferences we get

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

Most likely inspired by this form, Ross imposes the **third assumption** of **transition independence** on the kernel:

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

More general preferences than additive time-separable utility satisfy this form such as Epstein-Zin recursive preferences. At this point, we should pause to reflect how much this assumption has bought us. If we look at the LHS, in discrete time and state space world with m states, it will be a matrix with up to mxm degrees of freedom whereas the RHS just has m+1 degrees of freedom (h takes on m different values, and δ is also unknown). Thus, we reduced our search from m^2 to $m+1$.

From (9) and (10) above, we obtain an expression for state prices

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

In what follows, we **specialize our discussion to a discrete state space model**. Let

Ross Recovery Summary

By Edward Mehrez

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

where $U_i' \equiv U'\big(c(\theta_i)\big)$ or more generally, U' is any positive function of the state. Rewriting (11) in matrix form, we get

$$DP = \delta FD, \qquad (15)$$

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, φ_j / δ, on the diagonal,

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

We can then express the natural transition matrix F as

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

In addition, since F is a transition matrix, its rows must sum to one, we have the additional set of m constraints:

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

Rearranging, we get the characteristic root problem:

$$Pz = \delta z, \qquad (20)$$

where

$$z \equiv D^{-1}e. \qquad (21)$$

Finally we make our **fourth assumption** that P is irreducible which allows us to use the **Perron-Frobenius Theorem**. One of the results of the theorem is that all non-negative irreducible matrices have a unique positive characteristic vector, z and associated eigenvalue λ which corresponds to our lambda. This essentially gives us the main theorem:

Ross Recovery Summary

By Edward Mehrez

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.*

From this, we have that δ is the maximum characteristic root of the price transition matrix, P. Furthermore, another result of Perron-Frobenius gives us that this root is bounded above by the maximum row sum of P, which are the interest rate factors.

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

Furthermore, if the riskless rate is the same in all states, we get the surprising result

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.*

Multinomial Recovery (§ III)

Here Ross extends the theorem to an infinite horizon multinomial Lucus tree setting. It's useful to note that the three of the four sufficient conditions discussed hold in this setting. First, it is still a Markov process. Second, transition independence is directly assumed from assuming time additive utility set-up so that state prices will have the form of (10). Third, irreducibility still holds as any state almost surely is revisited in finite time. The assumption which is not met is the time independence assumption. The payoffs of the tree grow (or shrink) with time so it is another state variable. Under these assumptions we get Theorem 4:

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.*

An Example, Comments, and Extensions (§ IV)

In this section, Ross goes on to demonstrate the recovery theory in two different ways: first, he shows for a "static" example, that given both the utility function (CRRA in his example) and the stock price distribution (lognormal) that using the recovery theorem, one can recover the natural probability measure using the pricing kernel (SDF) and the state prices. In the given example, the utility function and the price distribution are used to derive the expressions for the SDF and the state prices (through the Black-Scholes-Merton formula). He then calibrates the standard deviation and mean return of the price process (price distribution) to market data and shows that the resulting recovered transition probability distribution coincides with the lognormal distribution. This section (Section IV) proves to be deeper than just a verification of the theorem for an example. First, Ross points out that although the theorem was proven for the discrete and multinomial cases, it still seems to recover with a continuous distribution when considering a static problem (moving from one known initial state to an unknown state). However, once the dynamic problem is considered where one first transitions from a known state to an unknown state, then from the unknown state to another unknown state (from time 0 information set perspective), problems arise in that no implicit truncation of the distribution can be used (see pg 632-633).

Applying the Recovery Theorem (§ V)

In this section, Ross relying on a rich market for European options to numerically approximate the state prices. To do so, we first note that a European call price with strike K and time of maturity T can be expressed as

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

From Breeden and Litzenberger (1978), and intuitively from naively applying the FTC, we have

$$\mathrm{p}(\mathrm{K},\mathrm{T}) = \mathrm{C}_{\mathrm{kk}}(\mathrm{K},\mathrm{T}) \qquad (83)$$

The RHS of this is approximated numerically from the second differences of the observed call prices. To apply the Recovery Theorem, we first have to estimate the mxm state price transition matrix.

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

Now, to index the states, we think of there being m possible states at any time 0 and time T – there is no growth in this formulation as in the multinomial model. Ross's notation in this section is pretty terrible, but for the sake of comparability with the paper I will maintain it and do my best to explain it. First, we denote each row of the transition matrix from time 0 to T, p^T, by $p^T(c)$ where c denotes the current state at time 0.

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

The arguments (k,T) denote the state at time T—the current state notation is suppressed in the vector list. This same notation is used in (83). Next, we denote P as the one period transition matrix. . One of the rows of P we get by using (83) to find the entries of $p^1(c)$ (cf. (85)). To find the rest of the entries of P, we use the following Markovian relationship

Ross Recovery Summary

By Edward Mehrez

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

where m is the number of states. This is a system of m^2 equations in the m^2 variables P_{ij}, and since we know current state prices, we can solve this.

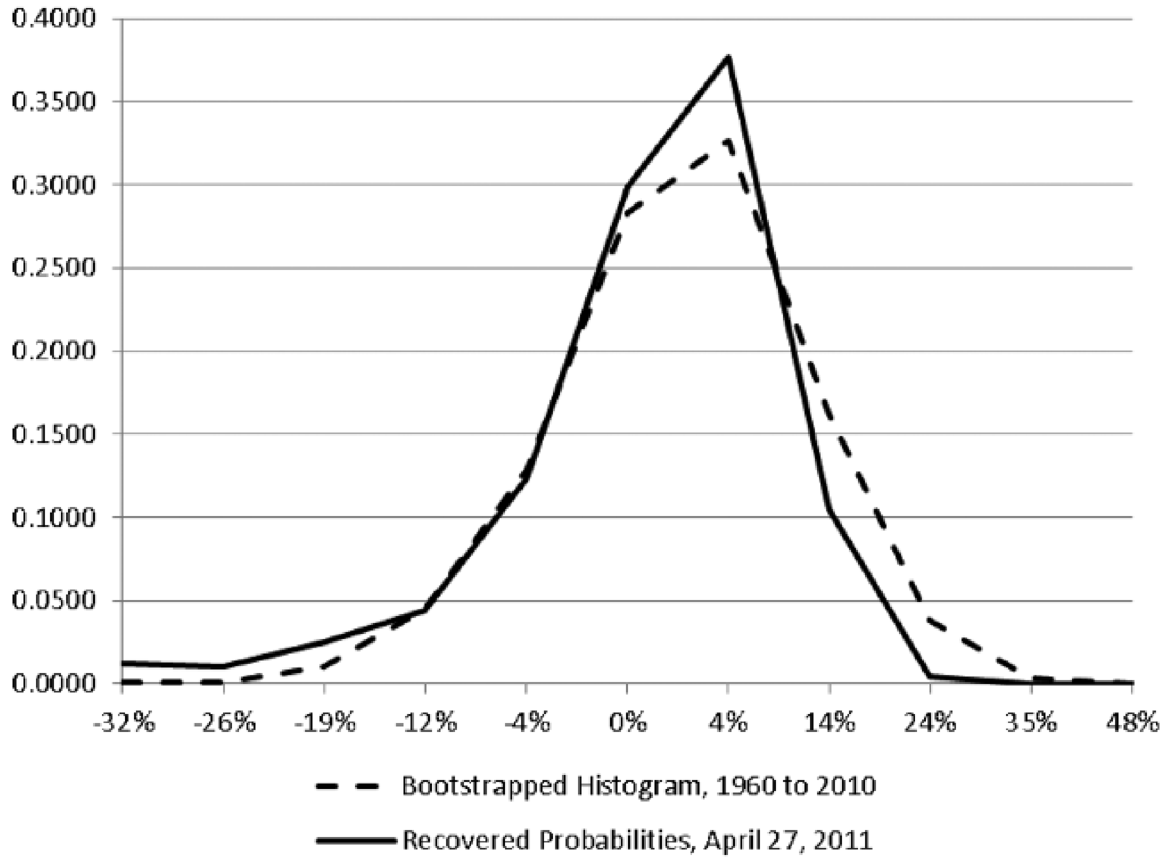

Figure 2. The recovered and the bootstrapped natural densities.

Figure 2 show the recovered densities vs the bootstrapped ones. Ross points out that the recovered density has a fatter left tail and suggests that this provides support to the recovery method as we should expect the subjective density to have a fatter left tail (fear of disaster) than the actual probability density.

Testing the Efficient Market Hypothesis (§ VI)

In Ross (2005), an alternative test to testing the EMH by finding an upper bound to the volatility of the pricing kernel is proposed. The recovery method allows us to find a number for this upper bound. In particular, from the Hansen-Jagannathan bound there is a lower bound on the volatility of the pricing kernel, ϕ:

Ross Recovery Summary

By Edward Mehrez

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

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.

Recovery gives us an estimate $\sigma^2(\phi) = 0.1065$. Next, we can decompose excess returns, x_t, on an asset or portfolio strategy as (see Ross(2005))

$$x_t = \mu\left(I_t\right) + \varepsilon_t, \qquad (90)$$

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

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

Rearranging and recalling (87) yields an upper bound on the R^2 of the regression:

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

Given the estimate of $\sigma^2(\phi)$ and interest rates at 0, this means that 10% of the annual variability of an asset (or portfolio) return is the maximum amount that can be attributed to movements in the pricing kernel with 90% idiosyncratic in an efficient market.