An Overview of the Recovery Theorem

Dhruv Madeka

Bloomberg LP

October 18, 2016

Overview

Introduction

- Historically, estimating physical probabilities for stock prices has been a statistical affair
- No concept of implied market distributions, a parametric class is estimated by using historical observations along with statistical or econometric methods (MM, GMM etc.)
- But payoffs extend into the future, so logically prices **should** contain embedded information about the market's view

Risk-Neutral Pricing

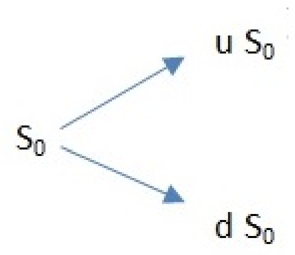

- We do not know if the flip is symmetric. In fact, we do not **need** to know for the pricing rule we are about to describe
- The prices depend on the state of the world rather than the probabilities of achieving these states

Risk-Neutral Pricing

- The familiar construction of choosing portfolios that pay 1 in state H and 0 otherwise and vice-versa provides the surprising result:

$$\Phi_H^* = \left(\frac{1}{1+r}\right)\frac{(1+r)-d}{u-d}$$
$$\Phi_T^* = \left(\frac{1}{1+r}\right)\frac{u-(1+r)}{u-d}$$

Continuous Time

- If the stock-price follows an Itô-Process under the real-world measure $\mathbb{P}$:

$$dS_t = \mu(t, S_t)dt + \sigma(t, S_t)dW_t \quad (1)$$

- We can construct the Radon-Nikodym derivative:

$$\eta_{t,T} = e^{-\frac{1}{2}\int_t^T \theta_s^2 ds - \int_t^T \theta_s dW_s} \quad (2)$$

where $\theta_t = \frac{\mu_t - r_t}{\sigma_t}$

- Construction of risk-neutral measure eliminates mean from the picture:

$$dS_t = r(t, S_t)dt + \sigma(t, S_t)dW_t^{\mathbb{Q}} \quad (3)$$

Black-Scholes Merton

- Black-Scholes PDE:

$$\frac{\partial f}{\partial t} + rx\frac{\partial f}{\partial x} + \frac{1}{2}\sigma^2 x^2 \frac{\partial^2 f}{\partial x^2} = rf \tag{4}$$

- It has been fortunate that the growth rate μ has not (yet!) entered the picture
- Minimum-Variance Unbiased Estimate:

$$\hat{\mu}_t = \frac{1}{t}\ln\frac{S_t}{S_0} + \frac{1}{2}\sigma^2 = \mu + \frac{\sigma W_t}{t}$$

$$\implies var(\hat{\mu}_t) = \frac{\sigma^2}{t}$$

$$\implies \text{If } \sigma = 20\% \text{ then } |\mu - \hat{\mu}_t| < 1\% \iff t > 1521 \text{ days}$$

The Recovery Theorem

- In 2013, Ross presented a theorem which identified the conditions under which a unique physical distribution for market indexes could be recovered from option prices:
- The Recovery Theorem states that:
  "If there is no arbitrage, 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."

The Recovery Theorem

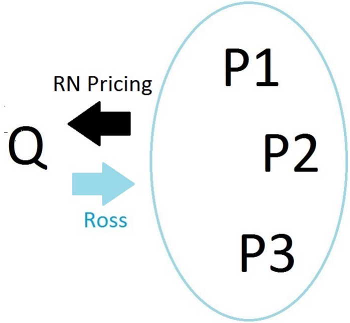

Is it really **that** surprising?

- It has been well understood for a while that it is incorrect to assume that the true measure's role lies only in providing a starting point for the risk-neutral measure
- The real connection between derivatives and the true measure is a little more subtle and revealing
- Grundy (1991) shows how the non-central moments of the true distribution connect to the upper bounds of call option prices
- Lo and Wang (1995) construct a simple example where the trending OU process and GBM result in the same pricing formula, resulting in the drift affecting the diffusion for the underlying process

Fundamental Theorem of Finance

Dybvig and Ross (1987) state that:

- The following three things are equivalent:
  - No Arbitrage
  - Existence of a positive linear pricing rule for all assets
  - The existence of a (finite) optimal demand for an investor who prefers more to less

Stochastic Discount Factor

- Fair-Price P_t of any T-payoff Ψ_T:

$$P_t = \mathbb{E}_t[\underbrace{e^{-\int_t^T r_s ds - \frac{1}{2}\int_t^T \theta_s^2 ds - \int_t^T \theta_s dW_s}}_{\text{Stochastic Discount Factor } \xi_{t,T}} \Psi_T] \tag{5}$$

- Single variable that describes the entire risk-return trade-off of the market in equilibrium
- Has hedging and replication implications in derivatives pricing theory
- Economic connection to marginal rate of substitution exploited by Ross for derivation of the Recovery Theorem

Connection with Marginal Rates of Substitution

- Consider a two period model. An investor solves the problem:

$$\max_x u(c_t) + \mathbb{E}_t[\delta u(c_{t+1})]$$
$$\text{s.t. } c_t = w_t - x \cdot p_t$$
$$c_{t+1} = w_{t+1} + x \cdot p_{t+1}$$

- Solving for p_t yields the pricing formula:

$$p_t = \mathbb{E}_t[\delta \frac{u'(c_{t+1})}{u'(c_t)} p_{t+1}] \tag{6}$$

Setup

- Consider a two period, finite state world where $\theta \in \Omega$ denotes the possible states of nature, and where the underlying process is Markovian
- Consider an intertemporal model, with a representative agent who exhibits additively time-separable preferences and a constant discount factor δ
- Let $c(\theta)$ denote consumption as a function of state, then the representative agent solves the following optimization problem:

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

$$\text{s.t. } c(\theta_i) + \int c(\theta) p(\theta_i, \theta) d\theta = w$$

Ross' Derivation

- If we denote $\rho(\theta_i, \theta_j)$ to be the transition pricing kernel and $f(\theta_i, \theta_j)$ to be the natural distribution, then the price of a contingent claim is given by:

$$p_0 = \mathbb{E}[\rho(\theta)g(\theta)] \tag{7}$$

- Here, $\rho(\theta) = \frac{p(\theta)}{f(\theta)}$ and the risk-neutral distribution $\pi(\theta) = e^{(r_0 t)}p(\theta)$
- Ross begins with a modified version of (5) and (6) to obtain:

$$\frac{p(\theta_i, \theta_j)}{f(\theta_i, \theta_j)} = \delta \frac{U'(c(\theta_j))}{U'(c(\theta_i))} \tag{8}$$

Ross' Derivation

- Equation (8) is the equilibrium SDF solution for an economy with complete markets in which prices are defined by the FOC for the optimum and intertemporally additively separable utility
- Ross exploits the matrix form of these equations to obtain an eigenvalue-eigenvector problem
- If we define P to be the state-price transition kernel, F to be the natural transition kernel and D to be the matrix of (diagonal) marginal utilities

$$\text{Where } D = \frac{1}{U_1}\begin{pmatrix} U_1 & 0 & 0 & \\ 0 & U_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & U_N \end{pmatrix} \tag{9}$$

Overview of the Derivation

- Equation (8) can be expressed in matrix form as:

$$
\begin{aligned}
D \cdot P &= \delta F \cdot D \\
\implies F &= \frac{1}{\delta} D \cdot P \cdot D^{-1} \\
\text{But } F \cdot \overrightarrow{1} &= \overrightarrow{1} \\
\implies P \cdot \left(D^{-1} \cdot \overrightarrow{1}\right) &= \delta \left(D^{-1} \cdot \overrightarrow{1}\right) \\
\implies P \cdot z &= \delta z
\end{aligned}
$$

Unique Identification

- Ross exploits the Perron-Frobenius Theorem to recover a unique distribution
- The theorem states: An irreducible matrix has at most one positive eigenvector and the associated positive (real) eigenvalue dominates other eigenvalues (in absolute value.)

It's quite simple really!

- The Recovery Theorem is (in implementation) not much deeper than observing that we have a uniquely determined system of equations:

$$p(\theta_i, \theta_j) = \delta \frac{U'(c(\theta_j))}{U'(c(\theta_i))} f(\theta_i, \theta_j) \ \ldots \ m^2 \text{ equations}$$
$$F \cdot \overrightarrow{1} = \overrightarrow{1} \ \ldots \ m \text{ equations}$$

- Unknowns: $f(\theta_i, \theta_j), U'(\theta_j) \ \ldots \ m^2 + m$ unknowns

Carr and Yu (2012)

- Carr and Yu observed that the restriction Ross places on preferences are equivalent to a restriction on the dynamics of the numeraire portfolio
- They extend the theorem to single-dimensional time-homogeneous diffusions living on bounded intervals with regular boundaries at both ends, such as instantaneous reflection (the amusing $\cot(X_t)$ example)
- They observe that the infinitesmal generator of such diffusions is a regular Sturm-Liouville operator with a unique positive eigenfunction

Walden (2013)

- Walden studied the conditions for recovery for single-dimensional time-homogeneous diffusions living on an unbounded domain and found that it is possible if both boundaries are non-attracting
- A necessary and sufficient condition established by Walden is that:

$$\int_{-\infty}^{0} e^{-\int_0^x \frac{\mu(s)}{\frac{\sigma^2(s)}{2}} ds} dx = \infty$$

$$\int_{0}^{\infty} e^{-\int_0^x \frac{\mu(s)}{\frac{\sigma^2(s)}{2}} ds} dx = \infty$$

- Essentially Walden's condition is equivalent to having non-attracting boundaries (or a diffusion that doesn't go to ∞ very quickly.)

Qin and Linetsky

- In a complete extension, Qin and Linetsky derive the necessary conditions which ensure recovery for (conservative) Borel-Right Processes
- They show that under two conditions, unique recovery is possible for BRP's:
  - Transition Independence for the Pricing Kernel
  - Recurrence of the driving stochastic process
- BRP's cover most relevant financial models, including jump processes

Borel-Right Processes

- Recurrence of a BRP is defined as follows:
  Assume that the Lusin topological space E has at least 2 points. A Borel-Right Process is said to be recurrent if for each $B \in \mathcal{E}$*,* $R(x, B) = 0$ *or* $R(x, B) = \infty$*,* $\forall x \in E$*.*
- *Where* $R(x, B) \equiv \mathbb{E}_x[\eta_B] = \int_0^\infty P_t(x, B)dt$ *is the expectation of the occupation time (the potential measure)*
- *Basically, on average the process spends either a zero or infinite amount of time in every universally measurable subset of the state space E.*

Dupire's Comment

- Dupire noted that the assumption of time-homogeneity would mean that the Recovery Theorem interprets a declining term structure of volatility as a form of mean-reversion in price
- If a time-dependent Bachelier Model is fit to data generated by an OU Process, then the mean-reversion is interpreted as a declining term structure of volatility
- The converse interpretation would create an incorrect extraction of the true distribution for a time-heterogeneous data-generating process

Utility Based Criticism

- Huang and Shaliastovich (2013) derive an extension that incorporates recursive preferences
- Ross' assumption of an expected utility framework means that the model will overestimate the implied probabilities of bad states
- No preference parameter for the timing of the resolution of uncertainty
- Use of recursive preferences, such as Epstein-Zin etc. would provide a better idea of how the market times this resolution

Borovicka et al

- Assumption of stationarity neglects important martingale component in the multiplicative semimartingale decomposition of the stochastic discount factor
- Quoi?

Hansen-Scheinkman (2009)

- Hansen-Scheinkman extract from semigroup theory a decomposition result for Markovian pricing kernels
- If a pricing kernel is a positive semimartingale multiplicative functional then it admits the decomposition result:

$$\xi_t = e^{-\lambda t} \frac{\pi(X_0)}{\pi(X_t)} \hat{M}_t \tag{10}$$

- Here $\hat{M}_t$ is a martingale, $e^{-\lambda t}$ is the subjective discount factor and the ratio $\frac{\pi(X_0)}{\pi(X_t)}$ captures the rate at which payments are discounted at time t given the states at both times

Hansen et. al (2014)

- In a direct critique of Ross' work, Borovicka, Hansen et. al (2014) point out that the stationarity assumption on the stochastic discount factor is equivalent to assuming that $\hat{M}_t = 1$
- The argument remains that in the presence of this martingale component the recovered measure is actually the distorted long-run measure of Hansen-Scheinkman (2009) rather than the physical distribution
- They present macro-finance models where the martingale component is significant (non-separable preferences etc.,) and empirical evidence of its non-triviality

So what?

- Criticism aside, the empirical work of Audrino, Ludwig et. al (2014) shows that there is information contained in the recovered distributions
- They use neural nets and a regularization technique to obtain a stable transition matrix from marginals
- Trading strategies generated from recovered distributions of SPX Options show profitability
- Results match intuition with consistent positive equity risk premium and negative variance risk premium

Implications

- Observing the true distribution of implied returns provides a host of applications in risk-management, asset allocation etc.
- However, the time-homogeneity assumption seems to be both a requirement and severe drawback
- Recovery or pseudo-recovery for time-heterogeneous diffusions and non-Markovian processes needs to be understood
- Extensions should focus on empirical applications and tests of the assumptions

Questions?

The End