

The Recovery Theorem*

May 6, 2013

Q Group

Steve Ross

Franco Modigliani Professor of Financial Economics
MIT

Managing Partner
ROSS, JEFFREY & ANTLE LLC

*This talk is based on the based on 'The Recovery Theorem', forthcoming in the Journal of Finance

The Distribution of Returns

- The generic approach to estimating return distributions in financial markets is purely statistical
- Often we assume that the distribution of returns lies in some parametric class and use historical observations to estimate the unobserved parameters
- For example, we might assume that stock returns follow a diffusion process and estimate the historical risk premium
- To estimate the spread we can use historical standard deviations or use the current VIX
- To apply the results for financial decision making or for policy or on any forward looking basis we assume that future returns are drawn from the historically estimated distribution
- But, the payoffs of financial instruments extend into the future and surely their prices must be informative about the subjective distributions of the payoffs
- We do use this information in the fixed income markets

Predicting Interest Rates

- We use forward rates to tell us about anticipated future spot rates
- Forward rates aren't unbiased predictions of future spot rates but, rather, embed a risk premium as well
- To separate out the risk premium from the spot forecast we estimate a wide range of parametric interest rate models
- These models assume a particular form for the distribution of interest rates as well as for the risk premium and then use historical data to estimate the relevant parameters
- This is sometime augmented by using the current yield curve to fit the model or as an observation from the distribution to update the estimates of the unobserved parameters

Forecasting the Equity Markets

- For equities we don't even ask the market for a single point forecast like we do with forward interest rates
- Here is what we do in the equity markets:
 - We use historical market returns and the historical premium over risk-free returns to predict future returns
 - We build a model, e.g., a dividend/yield model to predict stock returns
 - We survey market participants and institutional peers
 - We use the martingale measure or some ad hoc adjustment to it as though it was the same as the natural probability distribution
- What we want to do in the equity markets - and the fixed income markets - is find the market's subjective distribution of future returns
- Our first step is to look at the derivatives market and option pricing models; these markets price a wide range of contingent future equity flows

The Binomial Model

- Arguably, the binomial model and its close cousin, Black-Scholes, are the most successful model in economics if not in all of the social sciences:



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

- The absence of arbitrage implies the existence of positive Arrow Debreu (AD) prices for pure contingent claims

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

- Notice that the sum of the prices is the discount factor and that $p(a)a + p(b)b = 1$

The Binomial Model

- The price, D , of any derivative or any asset with payoffs $D(a)$ and $D(b)$ is simply

$$D = p(a)D(a) + p(b)D(b) = \left(\frac{1}{1+r} \right) [\pi D(a) + (1 - \pi)D(b)]$$

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

- This is the risk neutral theory of no arbitrage where prices are the discounted expected value of payoffs under the risk neutral or martingale measure
- Equivalently, no arbitrage and spanning imply that there is a unique portfolio of the stock and the bond that replicates the payoff of the option and the cost of this portfolio is simply the price of the derivative
- Strikingly, the natural probability plays no role whatsoever in this theory for pricing derivatives
- Conversely, there is no way to use derivatives prices to find any information about the underlying natural distribution, f

Black-Scholes

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

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

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

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

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

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

where $N(\cdot)$ is the cumulative normal, K is the strike, and

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


Black-Scholes

- As expected, the Black Scholes solution is the discounted expected value of the payoff on a call under the risk neutral measure, i.e., under the assumption that the stock drift is the risk free rate
- Since the expected return on the stock, μ , doesn't enter the formula, option prices are independent of that component of the natural distribution of the stock
- This is both a blessing and a curse;
- It means that we can price derivatives without knowing the very difficult to measure parameters of the underlying process
- But, conversely, it renders the market for derivatives useless for determining the natural distribution
- The challenge, then, is to see if it is possible to build a different class of model that will have the ability to link option prices to the underlying parameters of the stock process while at the same time being consistent with risk neutral pricing

The Options Market

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

The Volatility Surface



Surface date: January 6, 2012

Implied Vols, Risk Aversion, and Probabilities

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

$$\text{Put price} = \text{Risk Aversion} \times \text{Probability of a Crash}$$

- But which is it - how much of a high price comes from high risk aversion and how much from a high probability of a crash?
- The binomial model and Black Scholes not only provide no answer to this question, they also suggest that no answer is possible and, to some extent, these theories and we, have lost the intuitions of insurance
- In fact, much is made of the distinction between the risk neutral measure we infer from option prices and the latent natural measure and research usually carefully and rightfully separates these different concepts
- Nonetheless, one often encounters a careless blurring of this distinction
- For example, risk aversion is sometimes ignored and the skew of the vol surface is casually interpreted as proof that tail probabilities are large

The Recovery Theorem: The Model

- A state is a description of the information we use to forecast the market, e.g., the current market level, last month's returns, current implied volatility
- The probability that the system moves from state i to state j in the next quarter is f_{ij}
 $F = [f_{ij}]$ - the natural probabilities embedded in market prices
- $P = [p_{ij}]$ is the matrix of contingent prices, i.e., the prices of securities that pay \$1 if the system is currently in state i and transits to state j in the coming month
- The contingent prices are the Arrow-Debreu prices conditional on the current state of nature and they are proportional to the risk neutral or martingale probabilities
- We can find P from market prices and what we want is to use P to find F
- The next slide describes the relation between the contingent prices (or the martingale probabilities), p_{ij} , and the natural probabilities, f_{ij}

The Kernel (Risk Aversion) and Contingent Prices

- The contingent price, p_{ij} , depends on both the probability that a transition from state i to state j will occur and on market risk aversion
- Since purchasing a contingent security protects against the consequences of state j it is a form of insurance and, like any insurance, contingent prices are a product of these two effects and time discounting:

Contingent price of \$1 in state j given we are now in state i = $\delta \phi(i,j) f_{ij}$

where δ is the market subjective discount rate and $\phi(i,j)$ is the pricing kernel, i.e., the marginal rate of substitution that captures risk aversion

- The absence of arbitrage implies the existence of both the risk neutral measure and a positive pricing kernel
- But this is a very different approach than that of Black Scholes and risk neutrality
- The following slides show how to find the discount rate and risk aversion and isolate the natural probabilities from the market's risk aversion

The Recovery Theorem

- Keep in mind that we observe P and we are trying to solve for F
- The pricing equation above has the form:

$$p_{ij} = \delta \phi_{ij} f_{ij}$$

- Knowing P there is no way to disentangle the kernel from the probabilities;
- There are $2m^2 + 1$ unknowns and only $m^2 + m$ equations
- Suppose, though, that the kernel has the form of the marginal rate of substitution of a representative agent with additively separable preferences

$$\phi_{ij} = U'(C(j))/U'(C(i)) = \phi(j)/\phi(i)$$

- Now we can separate ϕ from F - see Appendix

A Simple Binomial Example

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

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

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

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

and

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

- This is four equations in five unknowns, ϕ_1 , ϕ_2 , f_1 , f_2 , and δ which, with some manipulation reduces to a quadratic in δ with a unique positive solution
- Without state dependence we would have

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

two equations in the four unknowns with no unique positive solution

Recovery and the Binomial and Black Scholes Models

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

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

but not μ and φ separately

- Unlike the traditional models, the interest rate is state dependent so while we have spanning and no arbitrage we can no longer pin down the AD prices as a function of a single interest rate

Applying the Recovery Theorem - A Three Step Procedure

- Pick a date
- Step 1: use option prices to get pure securities prices
- Step 2: use the pure prices to find the contingent prices
- Step 3: apply the Recovery Theorem to determine the risk aversion - the pricing kernel - and the market's natural probabilities for equity returns as of that date

Step 1: Estimate Pure (Digital) Prices

- Options ‘complete the market’, i.e., from options, we can create a pure Arrow-Debreu security, a digital option that only pays off iff the market is between, say, 1350 and 1352, one year from now (see Ross and Breeden and Litzenberger)
- We can do this with bull butterfly spreads, made up of holding one call with a strike of 1350, one with a strike of 1352, and selling two calls with strikes of 1351
- The prices of these pure securities are $P(c,j,t)$ where t is the maturity of the security and c denotes the current state - these are the prices for a \$1 payoff in state j in t periods
- The following two slides illustrate these pure prices for maturities up to 3 years and for returns from 35% down to 54% up - this range corresponds to proportional movements of one standard deviation = 30% annually
 - Notice that the prices are lowest for large moves, higher for big down moves than for big up moves, and that after an initial fall prices tend to rise with tenor


Pure Security Prices for \$1 Contingent Payoffs

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

Priced using the SPX volatility surface from April 27, 2011

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

The Pure Security Price Surface




Priced using the SPX volatility surface from April 27, 2011

Step 2: Estimate Contingent Prices

- The prices we need are the contingent prices, p_{ij} for moving from state i to state j in, say, the next month
- The next slide illustrates how to get these contingent prices, p_{ij} , from the pure state prices
- The approach is an application of the forward equation for the conservation of probability:
- We find the contingent prices by moving through the table of pure security prices along possible market paths and different tenors and strikes

Contingent Prices and Market Paths



Contingent Forward Prices

Quarterly

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

Priced using the SPX volatility surface from April 27, 2011

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

Step 3: Apply the Recovery Theorem to P


- We can now apply the Recovery Theorem and solve for δ and the kernel, φ
- Given the kernel and the discount rate, δ we can solve for the natural probabilities

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

- The next slides compare these predictions for a given date with the historical


Risk Neutral Density

May 1, 2009




Implied Market Utility Function

May 1, 2009




Recovered Density May 1, 2009




Recovered Probabilities vs. Risk-Neutral Probabilities

May 1, 2009




Recovered Probabilities vs. Historical Probabilities

May 1, 2009



Recovered Cumulative vs. Historical Cumulative

May 1, 2009



The Bootstrapped (Historical) and Recovered Probabilities On May 1, 2009 for May 1, 2010

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

Recovered and Historical Statistics

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

Some Applications and a To Do List

- We have to test the recovery method by estimating the future market return probabilities in the past, e.g., daily or monthly, and then comparing those estimates with the actual future outcomes and with predictions using history up to that day
- We can also compare the recovered predictions with other economic and capital market factors to find potential hedge and/or leading/lagging indicator relationships
- This could prove valuable for asset allocation and a host of practical issues
- Extending the analysis to the fixed income markets is already underway by Peter Carr and his coauthors and in joint research by myself and Ian Martin of Stanford
- Publish a monthly report on the current recovered characteristics of the stock return distribution:
 - The forecast equity risk premium
 - The chance of a catastrophe or a boom

APPENDIX

Recovery Theorem Proof



The Recovery Theorem

- Letting D be the diagonal matrix with the m risk aversion coefficients, $\phi(j)$, on the diagonal, the basic pricing equation relating natural probabilities to contingent probabilities has the form:

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

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

$$F e = e$$

- Rearranging the basic pricing equation we have:

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

and since $F e = e$,

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

or

$$P x = \delta x$$

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

The Recovery Theorem

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

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