Predicting the Market

April 28, 2012

Annual Conference on General Equilibrium and its Applications

Steve Ross

Franco Modigliani Professor of Financial Economics
MIT

The Importance of Forecasting Equity Returns

- Estimating the equity risk premium to use for asset allocation, risk management, pension planning, actuarial analyses, and accounting – to mention just a few
- Forecasting market volatility, the probability of a crash – or a boom, and, more generally, we want to estimate the whole probability distribution of future returns
- We need the probability distribution of returns for risk control, e.g., VAR type computations and scenario analysis
- We could use the distribution to test if a strategy that generates an alpha in backtesting is likely it is to do so out of sample
- Because financial markets price securities with payoffs extending out in time, the hope that they can be used to forecast the future has long been a source of fascination
- This talk outlines a new technique for using the prices of equity derivatives to determine the market's subjective forecast of the probability distribution of future returns as well as the market's risk premium – the Recovery Theorem

Predicting Interest Rates

- In bond markets, we already use prices to tell us about the future
- Forward rates are rates we can 'lock in' today for lending or borrowing in the future
- At best, forward rates only tell us what the market thinks will happen and certainly not what will actually happen, but even that is of great use
- We form our own models and forecasts and we ask experienced market participants what they think
- But whatever their opinions or our views, we compare these forecasts with forward rates

Forward Rates

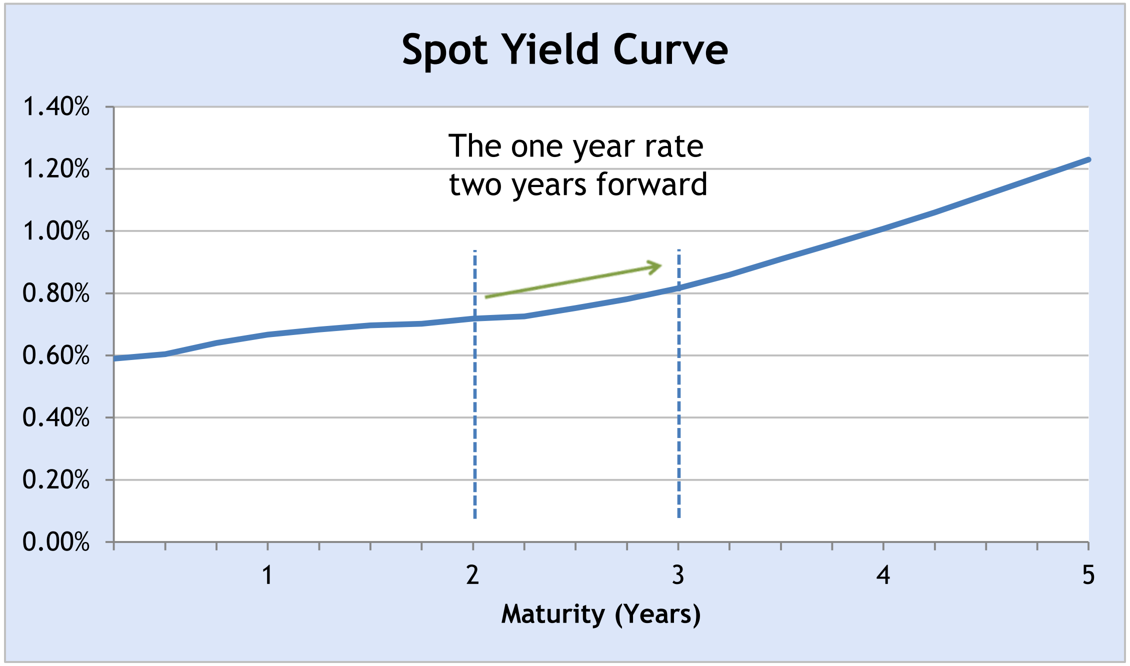

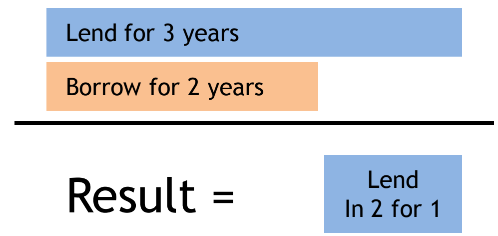

- E.g., the one year rate in two years is implied by the current spot two and three year rates
- An investor could 'lock in' lending at this rate (the in two for one) with spot rates by lending for three years at the three year rate and borrowing for two years at the two year rate (and the reverse is true for borrowing in two for one)

Predicting the Probability of Different Interest Rate Paths

- Forward rates are just one prediction out of the universe of possible future interest rates

- What we don't do with bond prices alone is answer a richer set of questions

- For example, what is the probability that the long rate will rise to 7% or higher in the next year, what is the chance that today's low rates will persist?

- To answer these questions, we need more than just forward rates derived from bond prices; we need contingent forward prices

- A contingent forward rate would be a rate for lending in the future that could be locked in today and that depends on what interest rates had done in the intervening period
  - Example: A rate we could lock in today for a one year investment to be made two years from now if the 10 year rate then is, say, 6%

The State of the Art in Forecasting Equity Markets

- For equities we don't even ask the market for a single forecast like we do with forward interest rates
- Here is what we do now 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
- We could then compare our forecasts and models of equity returns with the market's forecasts

The Options Market

- A put option on the market – the S&P 500 - is a security with a specified strike price and maturity, say one year, that pays the difference between the strike and the market iff the market is below the strike in one year

- A put option is like insurance against a market decline with a deductible; if the market drops by more than the strike, then the put option will pay 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

- We will use the rich (complete) market in equity derivatives, to derive contingent equity forward prices and then use them to unlock the market's forecast

Volatility and Implied Volatility

- The markets use the Black-Merton-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
  - The higher the implied vol, the more expensive are puts and calls
- Options with different times to maturity (tenors) and different strikes generally have different implied volatilities and the next slide displays this volatility surface
- Notice, though, that the market isn’t necessarily using the BMS formula to price options, only to quote their prices

The Volatility Surface

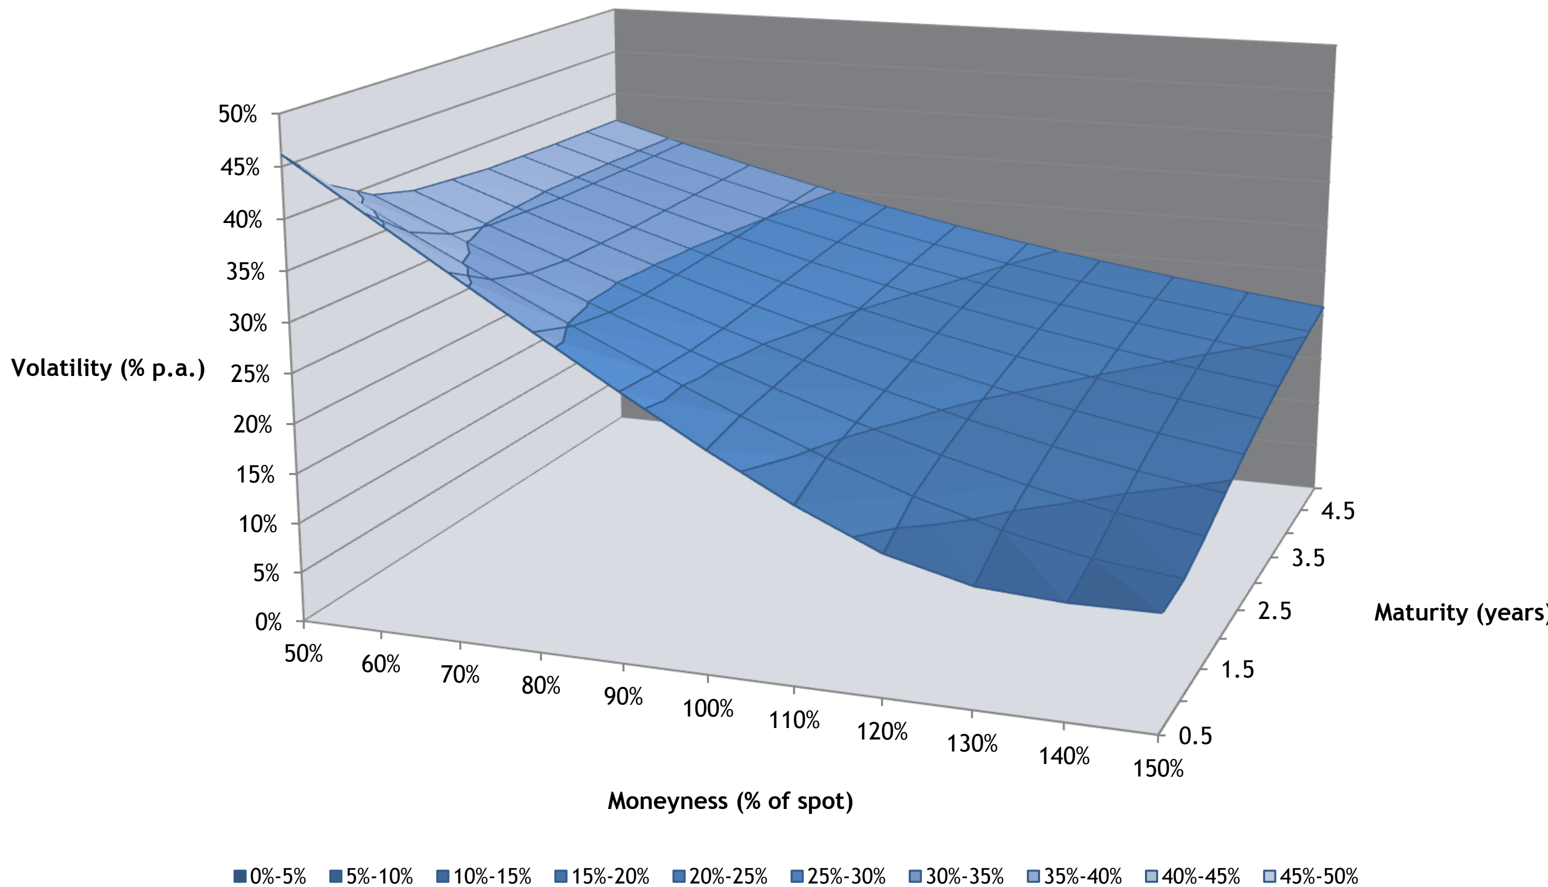

Surface date: January 6, 2012

Implied Vols, Risk Aversion, and Probabilities

- Implied vols for options on the S&P 500 are relatively high today (?) by historical standards, i.e., the level of the surface is high, particularly at the short end
- The vol surface is also highly skewed, i.e., prices for out-of-the-money puts are expensive because
  - The market thinks that the probability of a crash is historically high
  - And/or because the market is very risk averse
- Like any insurance, put prices are a product of these two effects:

Put price = Risk Aversion x Probability of a Crash

- But which is it – how much of the high price comes from high risk aversion and how much from a higher chance of a crash?
- The following slides develop a model that lets us separate these two effects and isolate the probability of a crash from the market's risk aversion

The Model

- A state is a description of what we use to forecast the equity 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 π_{ij} $\Pi = [\pi_{ij}]$ – the natural probabilities embedded in market prices

- $P = [p_{ij}]$ is the matrix of contingent prices, i.e.,, the prices of a security that pays $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 Π

- The next slide describes the relation between the contingent prices (or the martingale probabilities), p_{ij} , and the natural probabilities, π_{ij}

- The subsequent slides demonstrate how we can use this relation to find the natural probabilities, π_{ij} , from the contingent prices, p_{ij}

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:

$$\text{Contingent price} = \delta \times \text{Kernel} \times \text{Probability of a transition}$$

$$P_{ij} = \delta \times \text{Kernel} \times \pi_{ij}$$

where δ is the market's average discount rate

- From arbitrage alone we know that a positive kernel exists (and is unique in complete markets)

- The traditional approach to estimating this relationship is to either determine the kernel from other data or use historical data to estimate the transition probabilities

Estimating the Pricing Kernel I

- It is often said that we observe the natural distribution but not the martingale probabilities – let's see

- The graph on the next page shows the distribution of the martingale probabilities inferred from the vol surface, q(S)

- The following graph displays a lognormal estimated distribution of actual stock returns, f(S)

- The third graph displays the pricing kernel, i.e., q(S)/f(S)

- Clearly a ridiculously high $E[\phi^2]$ – not pretty

- Maybe its telling us that there are fat tails with high probabilities of catastrophes

The Martingale Measure

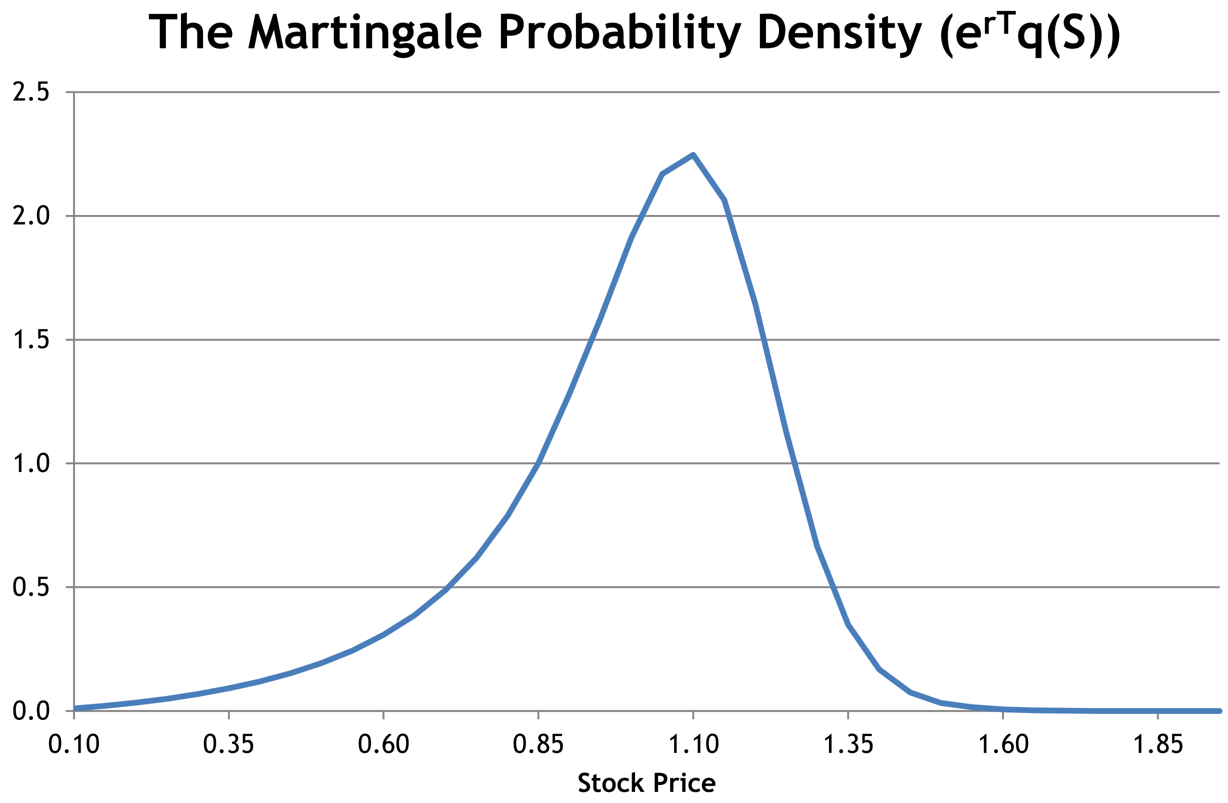

The Natural Measure

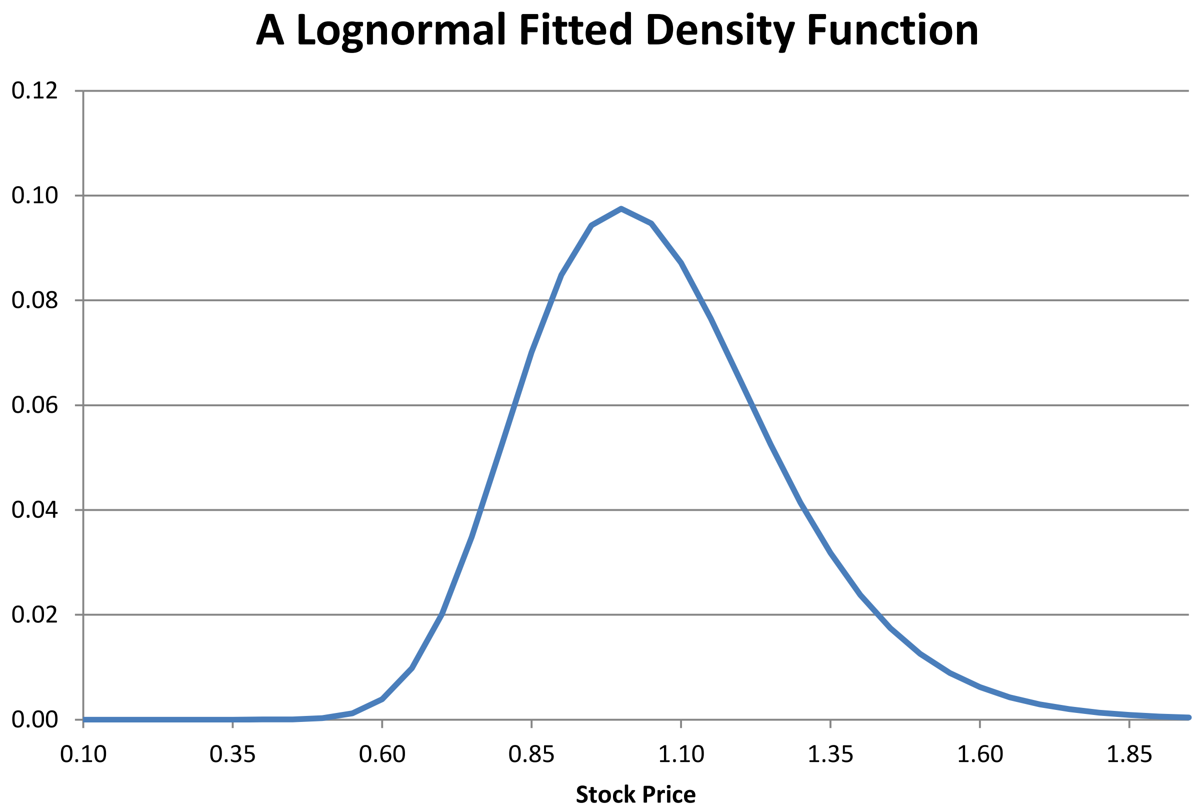

The Implied Pricing Kernel Density

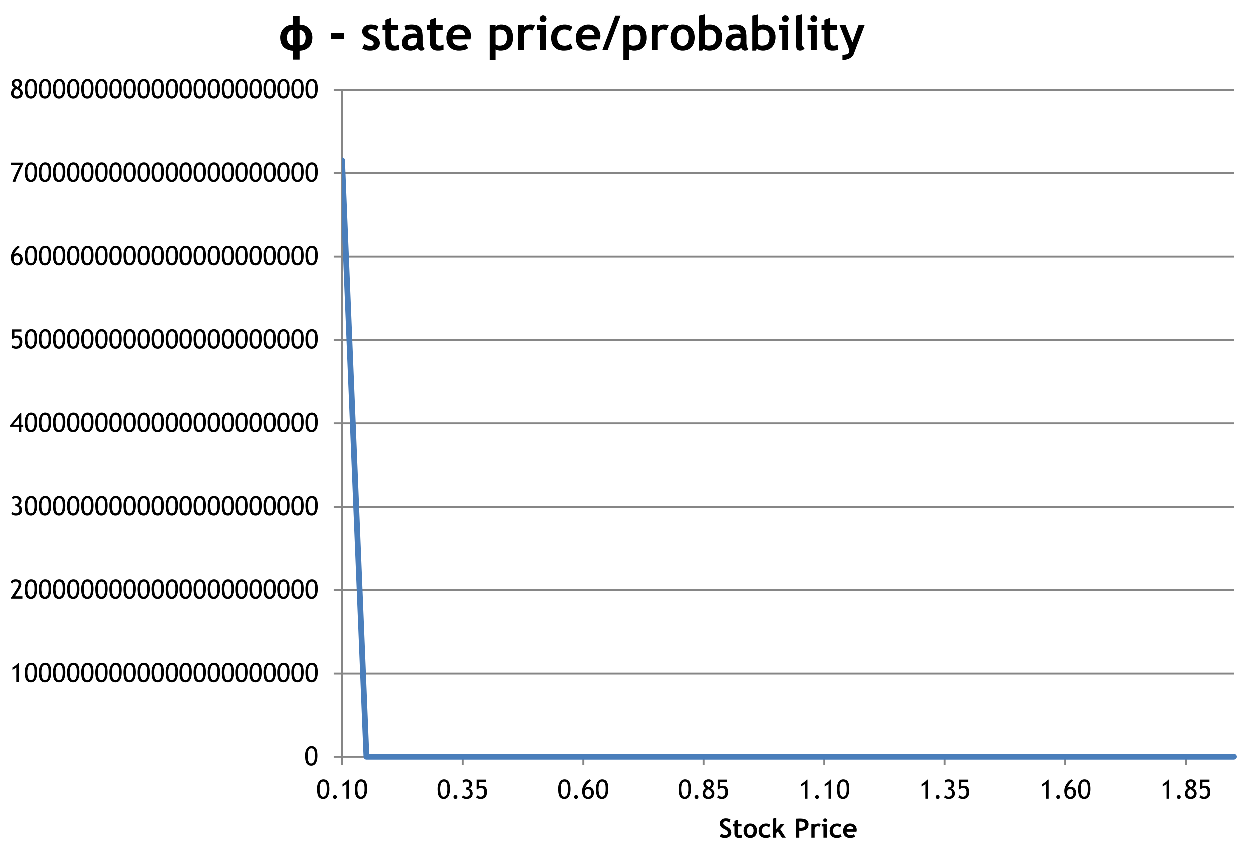

Estimating the Pricing Kernel II

- To get the natural density, suppose, instead, we use the histogram from bootstrapping stock return data (1960 to 2010)
- The graph on the next page shows the distribution of the resulting historical probabilities
- The next graph displays the pricing kernel, i.e., $q(S)/f(S)$
- Clearly we've learned our lesson and only plotted it where the probability is positive
- This may be appropriate because the state price density must be absolutely continuous with respect to the natural measure
- But it isn't really because there are positive prices where the natural measure is zero
- Again, its telling us that there are fat tails with high probabilities of catastrophes
- Interestingly, the resulting volatility estimate for $E[\phi^2] = 0.0058$ very tight

The Bootstrapped Natural Measure

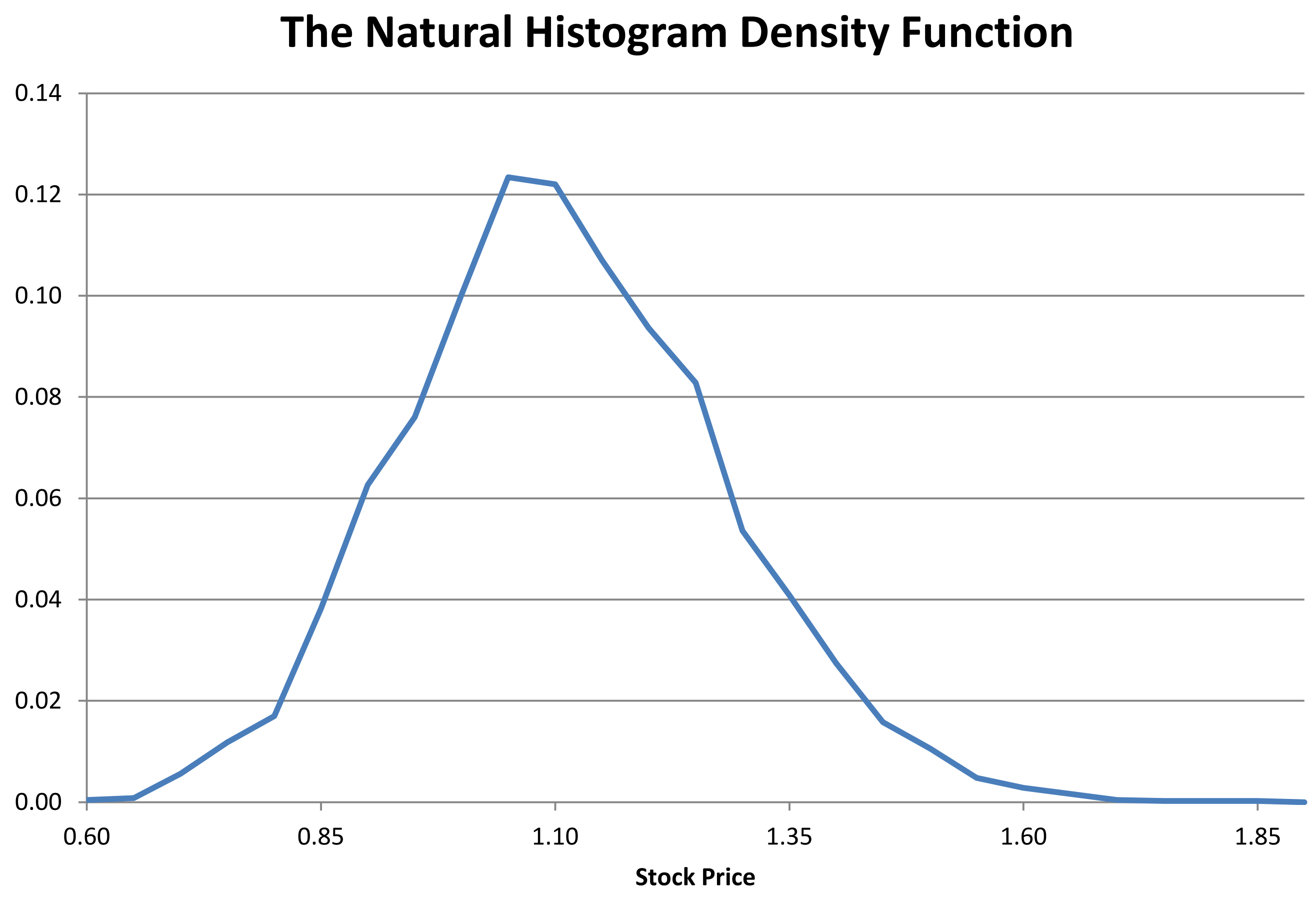

The Implied Pricing Kernel Density

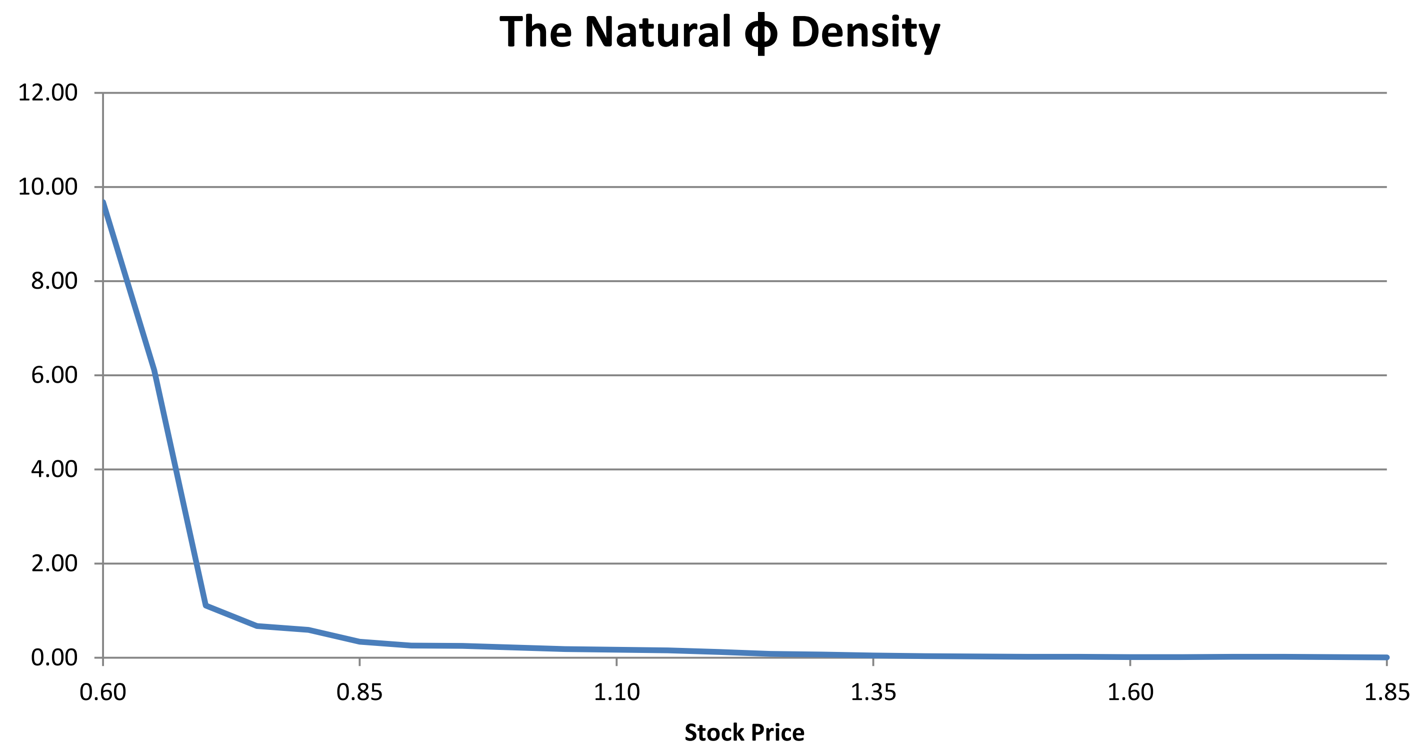

Estimating the Pricing Kernel III

- Let's try a different tack and use a stock distribution implied by a particular utility function

- Let the representative agent in this complete market have a constant coefficient of relative risk aversion, R, and set the kernel equal to the marginal utility, $\phi(S) = S^{-R}$

- This implies that the natural density, $f(S) = q(S)S^R$

- The next graphs display the kernels for three choices of R, R = 0, R = .5, and R = 5

- The last graph displays the resulting inferred natural density

- Not surprisingly, the lower risk aversion coefficients provide tight bounds on the volatility of the pricing kernel, $E[\phi^2] = 1$ and 1.03 for R = 0 and R = 1 respectively

- For R = 5, , $E[\phi^2] = 56$

- There is no obvious way to choose among these (insensible) estimates of the kernel

The Utility Natural Probability Density

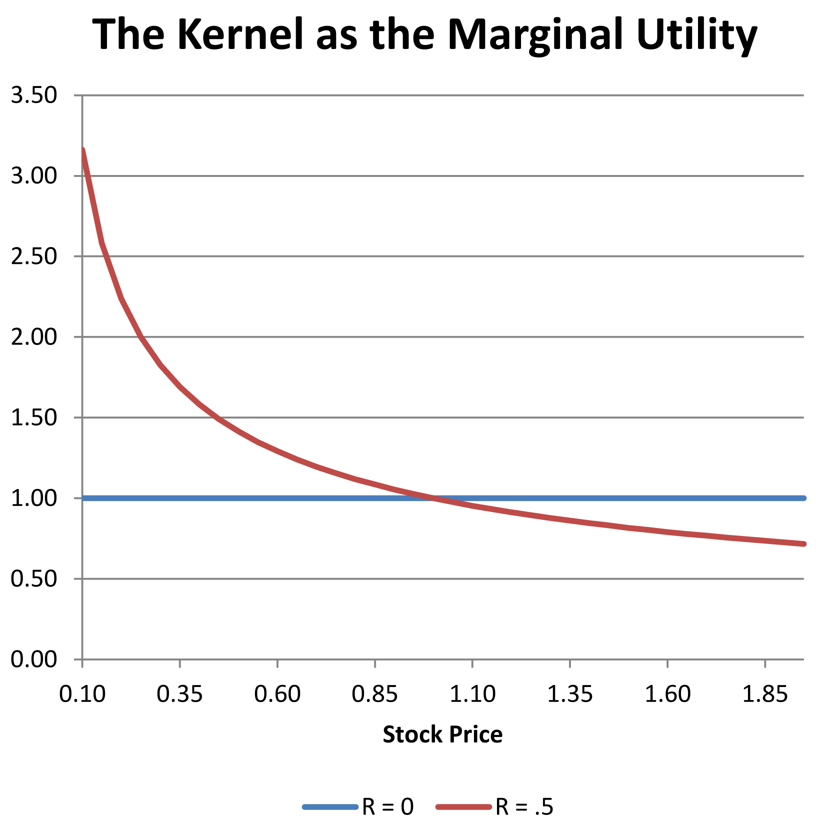

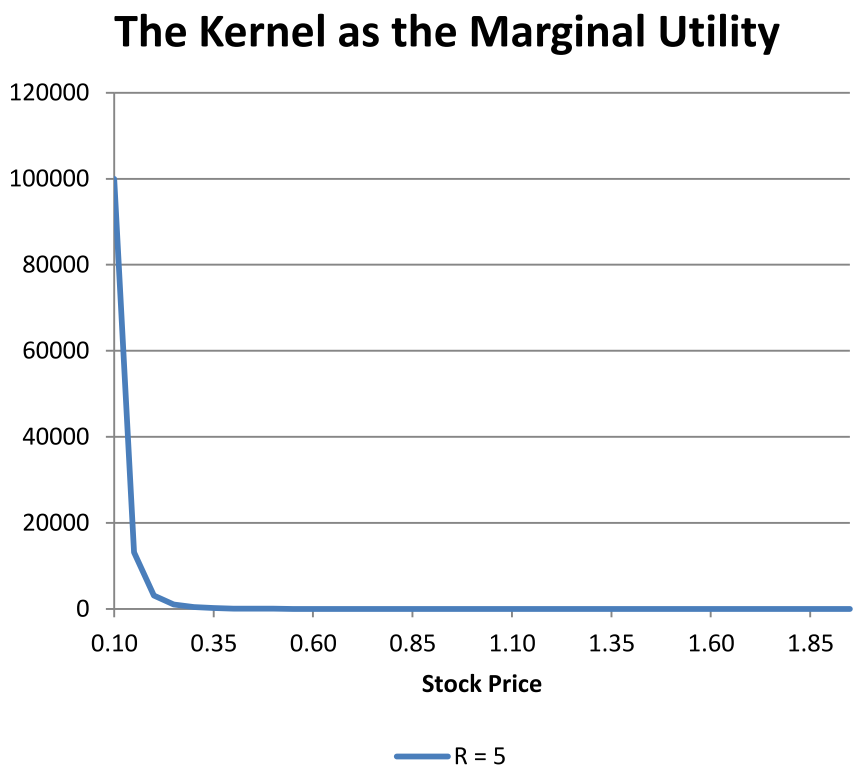

The Pricing Martingale Measure

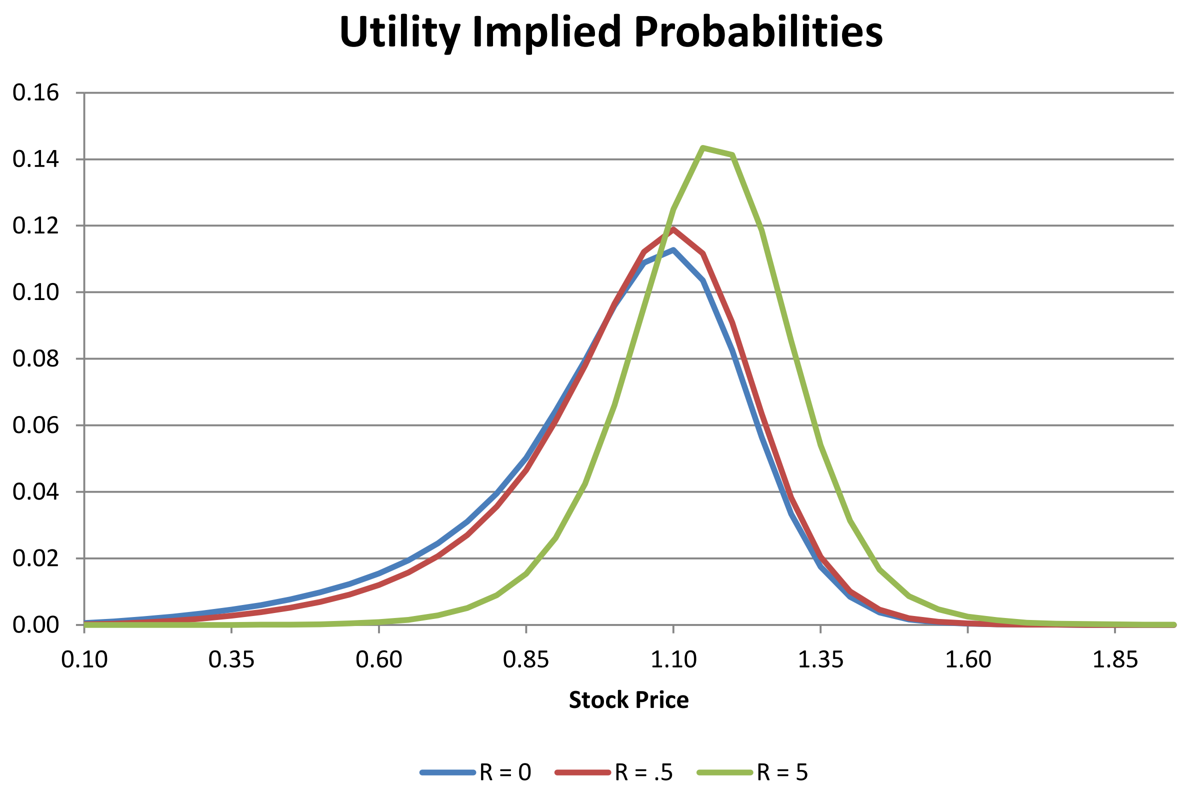

The Recovery Theorem

- Contrary to our initial intuition, both the kernel and the probabilities can be recovered from the price equation alone

- The pricing equation has the form:

$$p_{ij} = \delta \text{ x } RA \text{ x } \pi_{ij}$$

$$= \delta[U'(C(j))/U'(C(i))] \, \pi_{ij}$$

- A sufficient condition for this to hold is that the representative agent have an intertemporally additive separable utility function (or Epstein-Zinn recursive utility)

- What is critical is that the kernel is the ratio of the value of a function at the ending state, j, to the value at the initial state i, i.e., it doesn't depend on the transition from i to j

- If D is the diagonal matrix with the m risk aversion coefficients, RA(j), on the diagonal then the equilibrium equation has the form:

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

The Recovery Theorem

- Keep in mind that we observe P and we are trying to solve for Π
- Rearranging the basic pricing equation we have:

$$\Pi = (1/\delta)\, DPD^{-1}$$

and since the probabilities in the row sum add to one $\Pi e = e = \langle 1,\ldots,1\rangle$,

$$DPD^{-1}e = \delta e,$$

or

$$Px = \delta x$$

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

- This is a familiar eigenvector equation in mathematics and, assuming that P is irreducible (equivalently that Π is irreducible), from the Perron-Frobenius Theorem we know there is a unique positive eigenvector, x and eigenvalue δ
- From x we have D^{-1} and given δ we can now solve for the natural probabilities, Π

An Intuitive Argument for Recovery

- We have 11 possible price ranges (-35%, -29%, … , 54%) in our sample

- At each step the price transits from the current range – one of the 11 possibilities – to a new choice from the 11 possible price ranges

- This gives us a total of 11 x 11 possible price paths, and a total of 11 x 11 = 121 distinct equations, one for each possible path

  - These equations have 121 unknown probabilities and 11 unknown Risk Aversions

  - But since starting from any price, the 11 probabilities must add to one, we only need to find 10 probabilities for each starting price (as the 11^{th} will be implied); that is, 11 x 10 = 110 distinct probabilities

  - Thus, we have a total of 110 + 11 = 121 unknowns to find

- Since the number of equations equals the number of unknowns, we can solve for the the unknown probabilities and for the market Risk Aversions

Applying the Recovery Theorem – A Three Step Procedure

- I picked an arbitrary but interesting date, April 27, 2011

- 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: Estimating the Pure 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

Market Scenario	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
-35%	$0.005	$0.023	$0.038	$0.050	$0.058	$0.064	$0.068	$0.071	$0.073	$0.075	$0.076	$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	$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	$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
-8%	$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
19%	$0.007	$0.070	$0.129	$0.155	$0.166	$0.167	$0.164	$0.158	$0.152	$0.145	$0.137	$0.130
30%	$0.000	$0.002	$0.016	$0.036	$0.055	$0.072	$0.085	$0.094	$0.100	$0.103	$0.105	$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	$0.067
54%	$0.000	$0.000	$0.000	$0.000	$0.000	$0.000	$0.001	$0.001	$0.002	$0.002	$0.003	$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

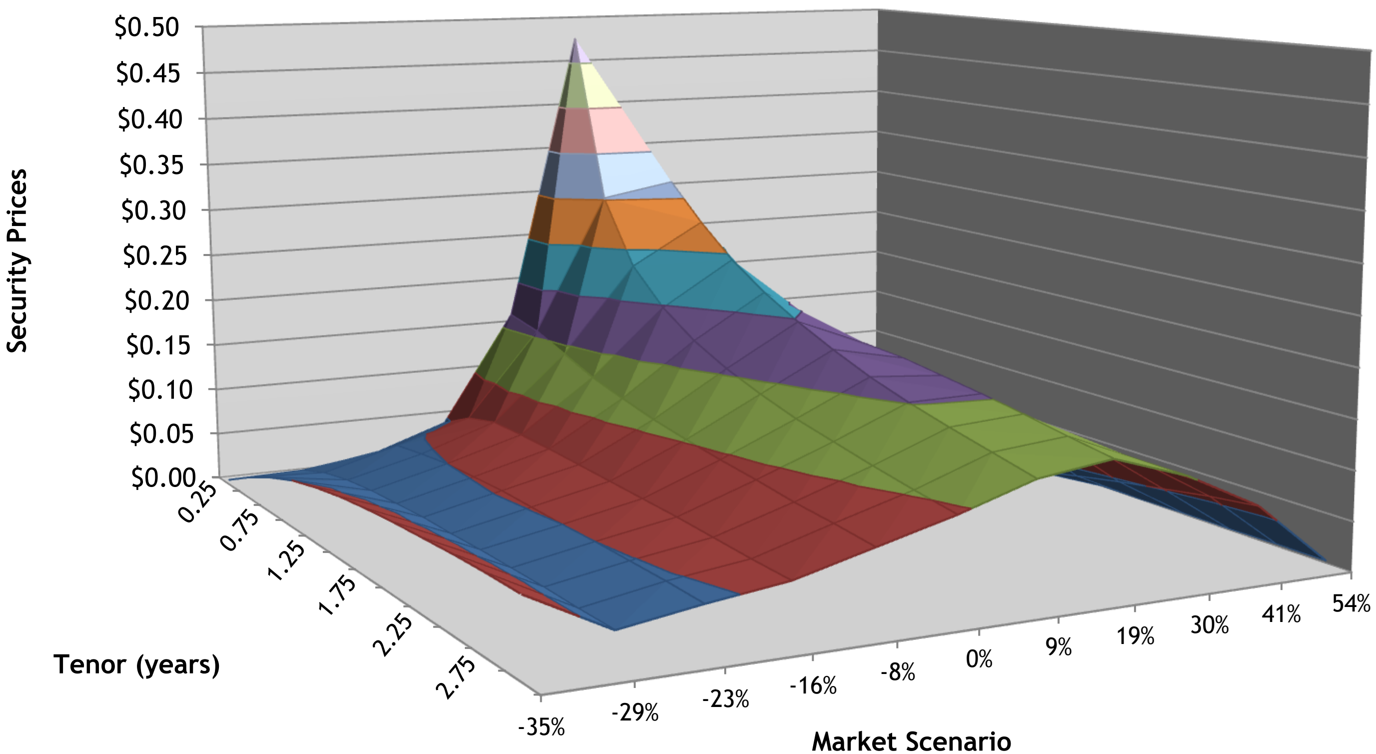

Priced using the SPX volatility surface from April 27, 2011

Step 2: Estimating the Contingent Prices

- For example, from the table p(c,+19% , 3) = $0.129
- Given the pure prices the next job is to get the contingent prices, p_{ij}
- The prices we need are the contingent prices, p_{ij} for, say, quarterly transitions
- As an example, the price of getting to up 19% by year end given that the market is down 8% by the end of the third quarter, is the contingent price of rising by three price range increments (9% each) in the last quarter
- We can find all these contingent forward prices by moving through the table of pure security prices along possible market paths and different tenors and strikes

Contingent Prices and Market Paths

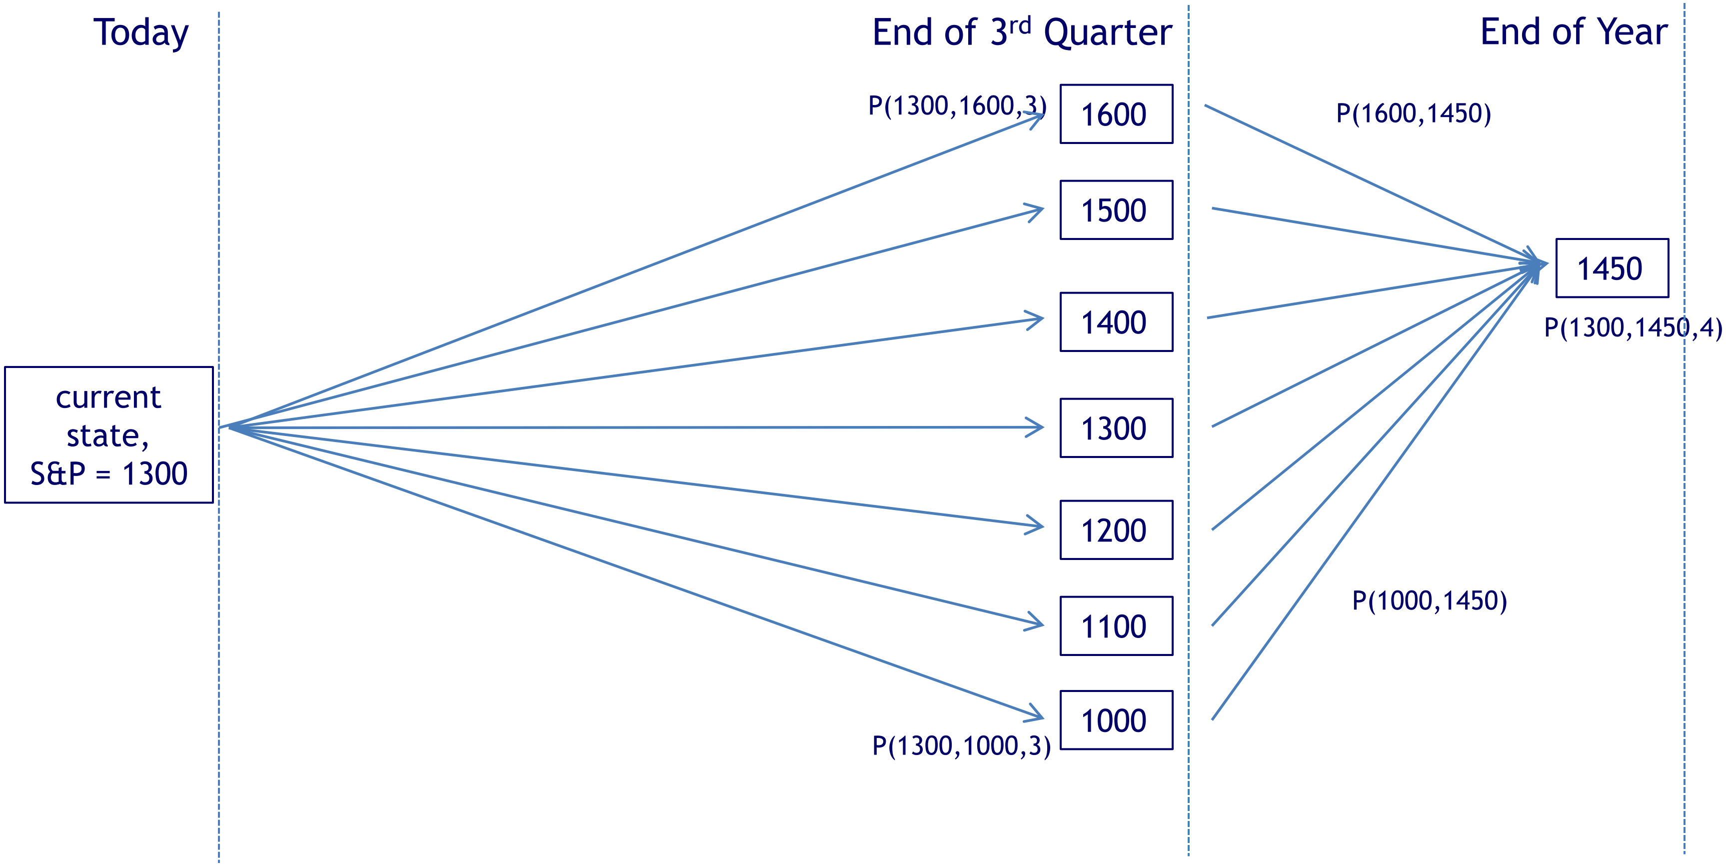

Contingent Forward Prices
Quarterly

Contingent Forward Prices

Market Scenario Final Period

		-35%	-29%	-23%	-16%	-8%	0%	9%	19%	30%	41%	54%
	-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
Market	-16%	$0.006	$0.044	$0.218	$0.390	$0.250	$0.057	$0.003	$0.000	$0.000	$0.000	$0.000
Scenario	-8%	$0.006	$0.007	$0.041	$0.211	$0.385	$0.249	$0.054	$0.002	$0.000	$0.000	$0.000
Initial	0%	$0.005	$0.007	$0.018	$0.045	$0.164	$0.478	$0.276	$0.007	$0.000	$0.000	$0.000
Period	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: Applying the Recovery Theorem to P

- We have 11 possible price ranges (-35%, -29%, … , 54%) in our sample
- At each step the price transits from the current range – one of the 11 possibilities – to a new choice from the 11 possible price ranges and we have found the (11x11) matrix of contingent forward prices, P
- We now apply the Recovery Theorem and solve for the eigenvalue, δ, and the eigenvector of P
- Dividing each price, p_{ij} by $\delta\,[U'(C(j))/U'(C(i))] = \pi_{ij}$, the natural transition probability
- The next slides compare these predictions with the historical numbers using the 60 years of stock market history from 1960 through 2010

The Bootstrapped (Historical) and Recovered Probabilities

Six Months Out

	Probabilities:		Cumulative Probabilities:	
Market Scenario	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

Calculated from the SPX volatility surface from April 27, 2011

- The historical probability of a drop of 32% or more was .08% and the probability of a decline in excess of 26% was .20% (using 5000 bootstrapped six month return periods)
- By contrast, the recovered probabilities of a catastrophic drop in excess of 32% is 1.2% and the probability of a decline in excess of 26% is 2.23%, over ten times higher

Recovered Natural Probabilities vs. Historical Probabilities

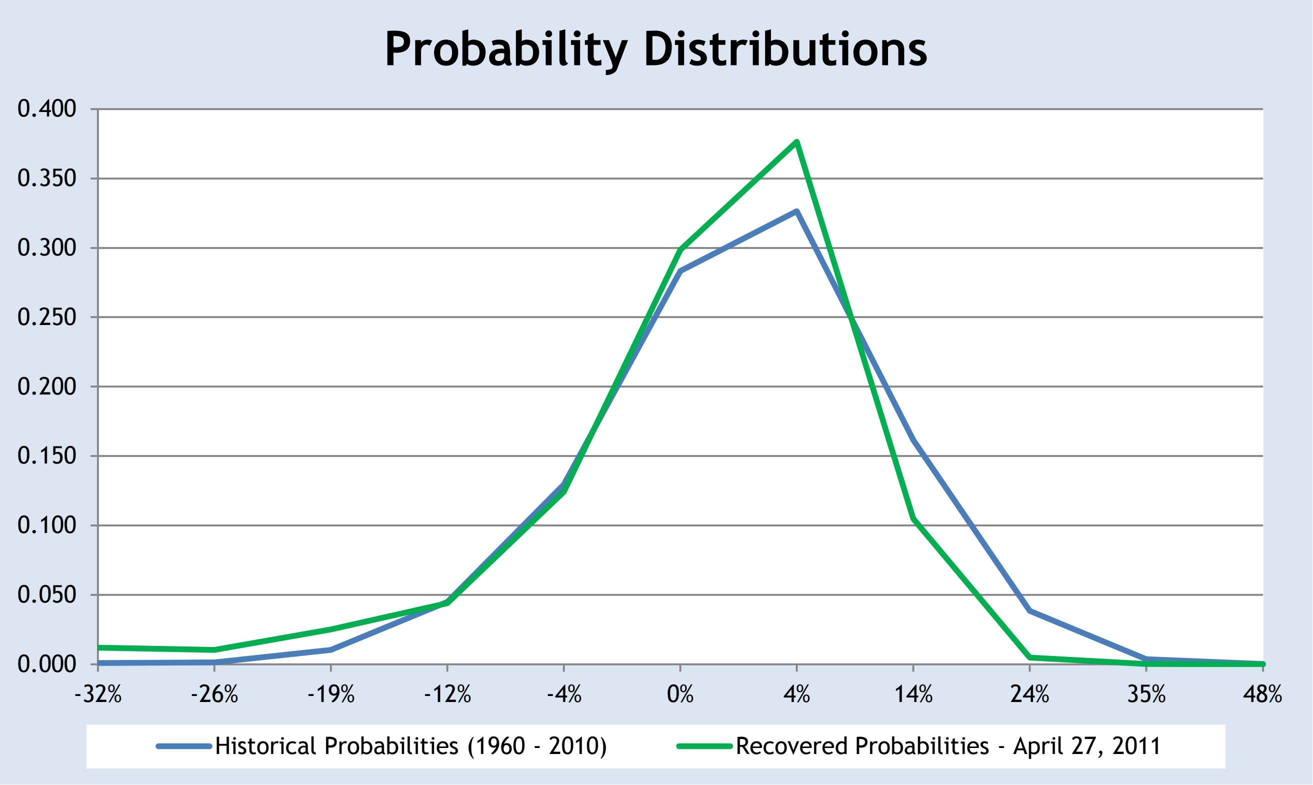

Annualized Recovered and Historical Statistics

Statistics					
	Recovered				Historical
Tenor	0.25	0.50	0.75	1.00	1.00
Mean	5.13%	5.52%	5.96%	6.16%	10.34%
Sigma	11.66%	14.04%	14.69%	15.00%	15.47%
Risk Premium	4.65%	5.28%	5.00%	4.31%	4.89%
Sharpe Ratio	0.399	0.376	0.34	0.287	0.316
ATM vol	14.53%	16.69%	17.71%	18.20%	

Calculated from the SPX volatility surface from April 27, 2011

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 will point us to new standards and metrics for managing both equity and equity derivative portfolios based on our more complete understanding of the implications of current prices
- Extending the analysis to the fixed income markets is already underway in joint research with 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