2001 Valuation Actuary Symposium
November 29–30, 2001
Lake Buena Vista, Florida

Session 5TS
Asset Modeling Concepts

Moderator: Michael J. Hambro
Panelist: Scott D. Houghton

Summary: Actuaries frequently include asset cash flows and asset valuations in strategic and regulatory projections of a company's business. In this session, panelists analyze the modeling considerations in projecting asset cash flows and asset valuations. Panelists also identify the risks associated with each asset class and the methods used to incorporate the risks into the projections. Topics include:

- *Interest rate and equity risk*
- *Setting discount rates for asset cash flows to obtain market values*
- *Determining the appropriate interest rate or equity scenario generator*
- *Uses of deterministic scenarios*
- *Prepayment and other risks for interest-contingent assets, such as mortgage-backed securities, asset-backed securities, and callable bonds*
- *Correlation between equity returns and interest rates*
- *Credit risk*
- *Types of credit-related risk*
- *Deterministic modeling of credit risk*
- *Stochastic modeling of credit risk*
- *Liquidity risk*
- *Sources of information for modeling*
- *Importing asset cash flows from another system versus modeling cash flows within the projection system*

Note: The chart(s) referred to in the text can be found at the end of the manuscript.

MR. MICHAEL J. HAMBRO: Actuaries project asset cash flow and value assets while performing several different functions, for example, regulatory, cash-flow testing, strategic asset/liability management, and corporate risk measurement and management. This session analyzes the modeling considerations in projecting asset cash flows and in valuing assets. A major focus of the presentation is on the risks that should be included, or at least considered, in asset projection and asset valuation. This is going to be especially pertinent for cash-flow testing in 2001. You have an environment where Treasury rates are very low and corporate spreads are very wide. Default rates are up to the highest levels they've been in at least a decade. You also have a very volatile mortgage-backed securities market.

Our presenter is Scott Houghton who is a consulting actuary at Aon Consulting in Avon, Connecticut. Scott specializes in asset modeling and insurance asset/liability risk management. He has extensive experience in pricing, modeling, and reserving for products that have embedded derivatives. Scott has also developed a commercial economic scenario generator for both interest rate and equity performance components. I am also a consulting actuary also at Aon Consulting in Avon, Connecticut. Scott has a lot of material to present, so I'm going to turn it over to him right now.

MR. SCOTT D. HOUGHTON: We're going to discuss various asset classes and their risks. We'll discuss how to capture these risks within our models, and we'll conclude with a brief discussion of other modeling considerations.

One of the main uses of our asset models is for regulatory projections. I'm sure all of you are familiar with cash-flow testing and the Actuarial Opinion and Memorandum Regulation. Another use is corporate modeling. This includes business planning and projections that are done to allocate capital and make sure we're using our company's capital efficiently. A final asset model use is strategic modeling. This category includes modeling that's done to determine the risk profile of assets and liabilities or to determine an investment strategy or interest crediting strategy.

What are the keys to successful asset modeling? I think the most important item is capturing investment risk. There are 13 bullet points that describe this session, and 10 of them are about different types of risk. Obviously, our data and assumptions are important. Our model construction, which is also important, is how we interpret and use the universe of information we have available in our model.

I'm going to talk about some asset classes and their risks. I'm going to start off with a very simple example. Let's discuss a Treasury bill quote, similar to what you might have seen in a paper back in August. The yields are a little bit lower now. It matures on November 8. The bid-and-ask interest rates are expressed on a dollar discount-basis. We usually see the yield to maturity in the quotations also.

The price of the bond with this yield would work out to about $99.47 for every $100 of par value. The cash flows are well defined, and the bill matures for its face value on the maturity date. Because the U.S. Treasury has the power to create money, there's no probability of default. The cash flows are certain. Given all these certainties, this asset is fairly easy for us to model. The real challenge for us in creating asset models is based on uncertainties. We'll look at those next.

I also have a Treasury note quote. It matures a little less than eight years from now in May of 2009. The cash flows are certain. There's this series of interest payments twice a year, in May and November, and there's a principal repayment in May of 2009 at the maturity. The price quote still has a convention in it. There's 32 per 100 of par. So the 102.27 means one hundred two dollars and twenty-seven thirty seconds, roughly $102.84 for every $100 of par value. There are bid-and-ask prices here. The bid price is the price at which the dealer will buy the bond from you. The higher ask price is the price at which the dealer will sell the bond.

The yield, 5.04, is much higher than on our Treasury bill, which was 3.49. The higher yield is there for a reason. It's compensation for risk. Most of the time, longer term bonds have higher yields than shorter term bonds. What if our lender or our insurance company needs the money earlier? What if we buy this bond and decide later that we don't want to tie up our money until

2009? The good news is there's a great secondary market for Treasuries. You buy at the asked price. The dealer will buy it back from you at the bid price. There's not a huge profit for the dealer in between these two prices. It's not like taking a new car off the lot. The other news, which could be good or bad, is that interest rates can change, and that affects the bid and the asked prices. To capture this in our model, we need to model future interest rates, so let's talk about the different ways of doing that.

The first type of interest rate model is a Treasury curve. Its rates and securities are issued by the Treasury. We graph maturity on the horizontal or the x-axis and the yield on the vertical or y-axis. The second type of model is the swap curve, which is a graph of the fixed rates payable on an interest rate swap. Companies are moving toward the swap curve as a basis for their modeling. It reflects a move towards a more international economy. Earlier this year, there was concern about too much Treasury debt being retired. There was more concern about that before the president's tax cut passed and we started getting refund checks in the mail. Most of the time, interest rates increase as term-to-maturity increases. We say the yield curve slopes upward. There has been some recent yield curves that have taken a funny shape, but earlier in 2001, about the time the president's tax cut passed, the Treasury curve started looking more normal shaped with a characteristic upwards flow.

The way we model interest rate risk is to change the future yield curves. The pattern of future yield curves is called an interest rate scenario. Chart 1 shows the Treasury and swap curves for the end of second quarter or as of June 30, 2001. As we learned in the general session, the short-term rates are about 1.5% less than they are on this chart. The long-term rates are approximately the same. The term-to-maturity is on the x-axis or the horizontal axis. The interest rate yield is on the y-axis. The swap curve is higher on the chart than the Treasury curve because the swap counterparties and the banks that lend at the London Interbank Offered Rate (LIBOR) index that determine the swap curve are not as creditworthy as the U.S. Treasury.

FROM THE FLOOR: Scott, please explain the difference between the Treasury curve and swap curve.

MR. HOUGHTON: It's the rate that is payable on an interest rate swap. If you were to pay a fixed rate on a notional amount of principal, you wouldn't necessarily need to have a bond or anything. It could be a notional amount of principal, the fixed rate that a company would pay in order to get a variable or floating or LIBOR rate back that varies based on the LIBOR index.

I mentioned that changing the future yield curves forms the interest rate scenario. There are two ways to do this. The first is called deterministic. As the name implies, these scenarios are predetermined. The pattern of future yield curves is determined in advance. The most common are the New York Seven scenarios that are used for cash-flow testing. Sometimes there's more than seven, depending on what jurisdiction you're in and what the starting yield curve looks like. Many companies have their own deterministic scenarios, and in Canada, the whole reserving process isn't formula driven like it is in the U.S. It's really a gross premium valuation, and there are certain required deterministic scenarios that the appointed actuary has to use.

A second type of scenario is called the stochastic scenario. Stochastic scenarios are generated with a random or a stochastic process. There are two main types of stochastic interest rate models. There are risk-neutral models, and there are realistic models. Realistic models are also called equilibrium models, and many of us use the term equilibrium models because calling them realistic models implies that risk-neutral models are unrealistic when, in fact, they're really just a different tool that's used for different purposes.

Risk-neutral scenarios are mainly used for pricing or determining the market value of assets. In general, I'd be talking about assets that are more complex than my two Treasury examples earlier. We'll talk about ways of determining the market value of assets a little later in the presentation, but risk-neutral scenarios give results consistent with the way an asset would be priced on Wall Street. The underlying assumption is that all investors are risk neutral. What does that mean? A risk-averse investor demands a higher return for a riskier investment. In a risk-neutral world, if there were two investments with different returns, all of the risk-neutral investors would buy the one with the higher return. There's no regard for risk. All this buying of the asset with a higher return would drive the price up and the expected return down until the returns would be equal in the two investments. In a risk-neutral world, all investments have the same return as a risk-free investment. In the real world, investors want a higher return for taking higher risk.

The reason we use risk-neutral models for pricing assets is because there's a large body of work that demonstrates the price of assets, which does not depend on the risk tolerance of the investor. Different investors have different risk tolerances, but asset prices reach an equilibrium anyhow. Being risk neutral is an extreme example of a risk tolerance, but the asset price doesn't depend on risk tolerance. It's much simpler to calculate the market value of an asset in a risk-neutral environment. The math is much easier in a risk-neutral world. We use risk-neutral models, not because the risk-neutral assumption is reasonable. We use them because it makes the math a whole lot easier, and it gives us the same answer anyway.

I'll give some examples of when I would use risk-neutral models. Say I'm evaluating a new security because I think it's appropriate for my company's investment portfolio. I want to know if the price the broker is offering me is a fair price or not. In another example, say I'm a reinsurer, and I want to reinsure a block of annuities and universal life. We'll say it's financial reinsurance. I'm going to assume the responsibility for paying the benefits on these liabilities. We'll say the mortality is predictable, so the mortality risk isn't material to me. We'll say the borrowers have options like surrendering to get a better deal elsewhere if interest rates rise. We'll assume I can model the behavior of the borrowers, and I can construct a portfolio of securities to replicate my obligations as a reinsurer. We'd say my company, being a reinsurer, might want to buy that portfolio to hedge our risk. We'd price the reinsurance quote that way. We'd charge a spread as an intermediary. That'd be another good example for using a risk-neutral model.

Realistic models are used to determine the range of results for a portfolio of assets and liabilities using historically based parameters. They include a risk premium for riskier investments. We would use a realistic model to model a typical life insurance portfolio. We'll say I have mortality claims that can fluctuate from year to year. I've modeled my mortality stochastically, and I want to find a risk profile of an asset portfolio that backs my liability portfolio.

Care must be used with realistic models to avoid exploitation of arbitrage opportunities that exist in the model but not in the real world. What does that mean? Most of the time our models will have some arbitrage opportunity. For example, by projecting interest rates at monthly or quarterly time steps, there'll be arbitrage opportunities by buying and selling securities that mature at different points in time within the month or within the quarter. That's not a big concern as long as we don't have model strategies that exploit that opportunity.

Say I have a stochastic generator that gives a five-year interest rate of 8% and a 10-year interest rate of 3%. I'll use simple interest so we can do the math in our head. What I would do is take out a 10-year loan at the 3%, 10-year rate. I'd invest the money in five-year bonds that are in the 8% five-year rate. Let's say I can borrow $1 million dollars for 10 years at 3%. I can invest it in the five-year bonds that earn 8%. At the end of five years, I'll have $1.4 million. That's five years at 8%, simple interest. That gives me the $400,000 plus the $1 million in principal. I'd stick $1.3 million in a vault since that's all I'd need to repay the 10-year loan at 3%. So I just made $100,000 without doing much work. That's a pretty blatant arbitrage opportunity. We'd want to avoid this in our models. It's possible to accidentally exploit something like that in our model strategies without realizing it.

I think it would be possible to devote the entire two-day meeting to choosing a stochastic generator. I'm going to stick with some pretty broad, general considerations. Session 9 also deals with this subject if you are interested in learning more. The type of modeling or the goal of the asset modeling is a big concern for us. Many models have both types of generators in them because we want to price assets, and we want to project future yield curves for modeling purposes. Cost and ease of use are always important considerations. Can I get the scenarios from the stochastic generator into the software that I use for asset/liability modeling (ALM) analysis?

I am often asked whether it is necessary to use stochastic scenarios? By just running deterministic scenarios, tell me what I want to know. The question I posed is, how well do deterministic scenarios capture interest rate risk? I'm going to use the New York Seven. How well do these seven scenarios capture interest rate risk? I'm going to try to answer that. We're going to set up a simplified asset/liability model. We're going to test it with the New York Seven scenarios, and we're going to test it with 1,000 stochastic scenarios. Then, we'll compare the results of both.

Let's look at a summary of the assumptions for the test I'm performing. There's a single premium deferred annuity product. There's $1 million of premium. The profits are retained like cash-flow testing. I keep all the profits and surplus that develop within the model. We're going to look at the ending market value of the surplus, which is the market value of the assets at the end of the projection. I'm going to subtract the cash surrender value of the liabilities at the end of the projection. I am also going to test the New York Seven scenarios. I'm also going to test 1,000 stochastic scenarios. I'm going to round the results a little bit.

As far as the single premium deferred annuity (SPDA) product goes, there's a 5% base lapse and there's a shock lapse after the surrender charge period ends. There are market-sensitive lapses. If there's a higher interest rate available, a policyholder can get higher interest rates someplace else, and the lapses are adjusted upward. Mortality is 70% of the Annuity Table A. There's an issue age of 55. There's a seven-year declining surrender charge, at 7% declining 1% a year. There are also some expense assumptions.

There's a $50,000 average size, which is used with the expense assumptions, and there are Commissioner's Annuity Reserve Valuation Method (CARVM) reserves. I'm using portfolio rate crediting with a 150-basis-point spread. The credited rate is reset once a year to the current portfolio rate, less the spread. I'm using a 35% tax rate with the deferred acquisition cost (DAC) tax capitalized.

I'd like to give you a feel for what this product looks like. When I actually run the test, I'm going to retain the profits, but I just did an initial test with the profits released. If I were to set up required capital equal to 200% of risk-based capital, I would need to sell $1 million dollars of this product. It is an initial capital infusion to set up the required capital and pay the acquisition expenses. What we'd see is that initial capital infusion and then a subsequent profit stream that the product generates. If I take the rate of return of this initial capital infusion and a subsequent profit stream, I get 13%, which is within the range that various companies would price an annuity product for.

I don't use the internal rate of return (IRR) as a profit result for my stochastic projections because some of the scenarios will generate funny profit patterns, and it's difficult to calculate an IRR. Sometimes there's no solution, and sometimes there's more than one. I'm using retained earnings, and my ending market value of surplus is my profit indicator. This should give a feel for what the product looked like.

As far as asset assumptions go, I'm using the scenario that was in place at the end of the second quarter. When there's excess cash flows, I'm going to invest half of them in 10-year corporate bonds and half in collateralized mortgage obligations (CMOs). The CMO is a middle sequential tranche of an agency CMO. There are prepayments that are appropriate for an agency CMO. When interest rates drop, the underlying mortgageholders will prepay the CMO, and my principal, as an investor, will come back early.

I'm deliberately taking some investment risks here in order to get a range of results. Resetting the credited rate annually on the annuity makes it a fairly low duration liability, but I'm buying 10-year bonds and longer term CMOs. So I definitely have a duration mismatch here. The CMOs have a negative convexity because of the call or prepayment provision. So I'm taking duration risk and convexity risk. I chose this investment strategy to get a range of results to compare my deterministic and stochastic scenarios, not because it's necessarily a good investment strategy for this SPDA product.

As far as the scenarios, I've used a realistic generator. There's a two-factor, lognormal interest rate model. There are short- and long-term interest rates, and there's a correlation between the changes in short- and long-term interest rates. There's mean reversion, which means that, over time, rates tend to revert to a mean or an average value. I've used parameters based on history for both the volatility and the mean reversion targets, and the average rates for mean reversion targets.

I create these scenarios shown in Chart 2. I input in all these assumptions, and I let a model run seven deterministic scenarios. I let the model run 1,000 stochastic scenarios, and I would get those results a few days later. These are the New York Seven scenarios. The scenario number is along the horizontal or the x-axis. The market value of surplus is on the y-axis. Scenario 1 in the New York Seven is the level scenario on the far left side. That gives me a market value surplus of $74,000 for an ending value. Scenario 2 is the gradually rising interest rate scenario. The result is negative. I have negative surplus at the end. If I did cash-flow testing with this, I would fail that scenario. Scenarios 3 and 4 are the other rising scenarios. They are rising and then falling, and there is an immediate increase of 3%. Five through seven are the falling scenarios, and those do a little bit better than the level scenario. This isn't a surprise for our result. As we know, our duration of assets was longer than our duration of liabilities. So we get good results when interest rates fall and poor ones when they rise.

Chart 3 shows the 1,000 stochastic results. I presented them here a little bit differently. The horizontal or the x-axis is now the surplus amount, and the vertical axis, or the y-axis, is the number of scenarios of the 1,000 that fall within a range. That's why I rounded the results to the nearest $5,000. I put the intervals of 5,000 on the horizontal axis, and the vertical axis represents the number of observations or the number of scenarios that fall into that range.

The graph gives us the distribution of results. It kind of looks like a normal curve that is skewed a little bit. You might notice that 23% of my results are below zero. If I was doing cash-flow testing, I'd fail 23% of the time. Like the New York Seven scenarios, this shows that perhaps it isn't a great investment strategy. The mean is 26,000, and the median is 31,000. The mean is a little less than the median because there is a big left tail in the distribution. I'd expect this graph to smooth out with more than 1,000 scenarios. One thousand scenarios do not represent a really large enough sample for this type of analysis. If we have any pricing actuaries in the audience, the level scenario result is 74,000, and that's the 93rd percentile of my curve here. That means that 93% of the time the scenarios come out worse than this, and 93% of the time, I don't get my pricing IRR of 13%. This is an extreme example. I've used that mismatch investment strategy, but 93% of the time, I don't get my pricing IRR. A level scenario is a very good outcome for an interest-sensitive product.

If I combine the results in Chart 4, the six squares can represent the results of the New York Seven scenarios. Some of you might ask, Why are there six squares if there are seven scenarios? There are three kinds of actuaries. Two of the results are very close together. So the fifth point (or the second point from the right) really represents two of the results. The line on the chart represents the stochastic scenarios. The same is true on the prior graph. The y-axis represents the number of observations, or the 1,000, and the x-axis represents surplus amount. For the New York Seven, the six squares, represent the surplus amount, and the height doesn't mean anything.

I have the three rising scenarios with the poor results on the left and the three falling scenarios with the good results on the right. What can we learn from this? I want to continue to focus on Chart 4. For this model and for these assumptions, the immediate increase of 3%, a pop-up scenario, is the worst case of the New York Seven. Ten percent of the stochastic scenarios are worse, shown in my left tail. Looking at these 10%, what could be worse than a 3% increase in interest rates for this product? There are two kinds of scenarios. There are a few that have an extreme increase in rates, and that's a few of them, but most of them reflect a case where interest rates increase, which we know is bad for this. Then, our assets mature. We purchase new ones in the model, and then interest rates increase again after that, which is very bad for this product.

In the New York Seven scenarios, the best result is the 98th percentile, and the worst result is the 10th percentile. The New York Seven scenarios, because they provide a fairly extreme stress test, give us an indication of where the tails of the distribution are. They give us a good indication of the range. About 88% of my stochastic results fall in between the best and the worst case of the New York Seven. They don't tell us the shape of the distribution, and they don't indicate the extremes.

I've talked about bonds and interest rate risk. I'm going to move on to our next asset class, which is stock. With stock or equities, there are two main types of risk. There's firm-specific risk, and there's market risk. Firm-specific risk can be diversified. We don't worry about it too much. It's not efficient to take that kind of risk. There's no reward for taking it. An exception might be an investment in an affiliate. Affiliate stock usually doesn't directly back our liabilities, but it could be considered if we're running a corporate model.

The second type of risk is market risk. The price or the market value of stock depends upon what happens in the stock market. We're going to talk about some equity models that capture that in a little while. First, we're going to talk about a related asset class.

Are equity separate accounts really subject to market risk? Separate accounts pass the investment risk to the policyholder, so why do we need to discuss this aspect. The policyholder is taking the risk instead of the company. Why do we care? Let's look at an example.

This is another modeled example. It's going to be similar to the variable annuity model we use to test deterministic scenarios. We don't have that bad investment strategy anymore that we used for the SPDA. We have equity separate account investments now. In the interest of time, I'm just going to look at some very high-level summary results. I'm going to vary the equity growth rate, whereas before I varied the interest rate. I'm going to vary the growth rate of my separate account. I'm also going to limit this to three deterministic cases of different equity growth rates. There's going to be a 150-basis-point mortality and expense (M&E) charge or asset-based charges instead of my 150-basis-point spread that I used before. I'm not going to hold the profits in the model. I'm going to distribute them so it looks like a pricing model. If my product generates profits every year, I'm going to distribute them. I'm going to take the internal rate of return of that profit stream and use that as my result.

I'm going to hold an approximation of 200% of risk-based capital, which is what companies typically might need to hold for a product like this. Any earnings that place my surplus position above that will be distributed and be part of that profit stream. I have some higher expenses that give me the results I want, but they also do reflect the marketing and the cost of supporting the separate accounts for variable products.

I run my test with three different equity growth rates, and I look at the results. I'm going to use the three different growth rates of 10%, 5%, and 0%. If I take the rate of return of that profit stream, my company's rate of return depends upon the growth rate of the separate account. If I look at the separate account growth rate, and I reduce it 5%, my company's rate of return goes down to about 5%. If I lower it 5%, my company's rate of return goes down to about 5% again. The reason I get these results is because my acquisition and my maintenance expenses don't scale down if the market does poorly. In addition, my main revenue item, which is my asset-based charges, scales down when the market does poorly, but my expense items don't.

How many people in the audience work for companies that write variable or separate account products? That's a good number. How many of you have policyholders that would call the 800 service number less often when the stock market declines? Nobody. Expenses, in fact, can actually increase as the market declines. You need people to answer all those calls and reassure your policyholders. What I've tried to show here is that if we have stocks in either our general account or our separate accounts, we have equity risk. How do we capture this equity risk? You have to use equity models. We use equity models for stocks, equity separate accounts, and equity-indexed products. With interest rates, there are deterministic and stochastic equity models, as well as risk-neutral and realistic. One type of model that's used quite a bit is the lognormal model, and we can use the lognormal model to either generate risk-neutral or realistic equity scenarios.

The most common lognormal model is called Brownian Motion shown in Equation 1. The variable p represents the stock price, or it could be the price of an index like the Standard & Poor's (S&P) 500 index. Delta p on the left is just the change in the stock price. That's equal to the initial or the starting stock price times the quantity in parentheses. Mu is the drift or the expected return. We want to use an expected return on my model with 10%. I just put in a 0.1

for the mu. Delta *t* is my time interval. Sigma is the volatility or just another word for the standard deviation. *z* is just a standard, normally distributed, random variable. So if I'm generating returns, I want to just take samples from a standard, normal distribution. That's what *z* reflects here.

$$\Delta P = P\,(\mu \Delta t + \sigma z \sqrt{\Delta t}\,) \qquad \text{[Equation 1]}$$

There's a theorem that's called Ito's lemma, which is in most financial modeling textbooks. One of the results of Ito's lemma shows that if a stock price follows Brownian Motion like this, then the change in the logarithm of the stock price over the time period is normally distributed, and a change in the stock price, which is delta *p,* has a lognormal distribution.

To create risk-neutral scenarios using a lognormal model, I use a drift rate of mu that's equal to the risk-free rate, which is the rate I could get on the Treasury security. To create realistic scenarios that include a risk premium, I use a drift rate or expected return that's higher than the risk-free rate. It's called a risk premium. There are other equity models besides the lognormal model. There are some more complex models that are beyond the scope of what I'm going to talk about during this session. Many of them have additional random variables besides the stock return. For example, the volatility, which is the standard deviation used to create the returns, can itself be a random variable. In the base lognormal model we assumed it to be constant.

One of the models currently being used quite a bit is regime switching. That's where the underlying stochastic model can change back and forth between two or more lognormal models that have different means and different standard deviations. The reason for using a model like that is an attempt to better match the actual distribution of stock returns.

Let's just take a look at Chart 5. It shows the distribution of stock market returns over the last 40-60 years. It shows the rationale for using regime-switching models that can better capture the tails of the distribution. It's more complex than an ordinary lognormal model. We'd want it to do a better job in order to justify using it. The solid line on my graph is the actual distribution of stock returns. The dotted line is a lognormal model with the same mean and standard deviation. The horizontal or the x-axis is the logarithm of the change in value of the S&P 500 index from one period to the next. For the small values, the logarithm of the value is approximately the return itself. The y-axis is just the number of observations in the S&P 500 that match this return on the x-axis.

The horizontal axis, or the x-axis, is the logarithm of the change in value. The dotted line is actually a normal distribution that I've plotted. It is not a lognormal line like it says. This is assuming that the distribution of the market returns is lognormal.

The actual data has some points out in the tails. There are some bumps in the solid line that aren't really captured by the lognormal distribution. A regime-switching model does a better job. The actual data has some periods of high volatility that are bunched together and they are better captured by a regime-switching model.

When we combine interest rates in equity models, it's fairly easy to do a deterministic scenario. If I have seven interest rate scenarios and eight equity scenarios, I can come up with 56 combinations. As far as stochastic scenarios go, if I start with 100 stochastic equity scenarios and 100 stochastic interest rate scenarios, and I combine those together, I end up with 10,000 combinations. Unless I have a real simple model, that's a lot to run. It's more efficient to create the interest rate and equity prices at the same time. In other words, I can run a smaller number of scenarios with variable interest rates and equity simultaneously.

In order to combine these, we want to consider the relationship between interest rates and stock prices. There's usually a negative correlation. If interest rates fall, the stock market goes up. If Alan Greenspan cuts interest rates unexpectedly, usually the stock market increases. The actual correlation does depend upon the historical time period.

We can look at some historical periods and the correlation between changes in interest rates and equity returns (see Table 1). There's a negative correlation, as we would expect. There's a period in time that roughly corresponds to when Ronald Reagan was president, where the data are positively correlated. It tends to muck things up by giving us a different answer when we start including it. My point here is that there's not a right answer as to what to use for a correlation. Many practitioners do give higher weight to the more recent data near the top.

TABLE 1

Historical Period Ending 6/30/01	Correlation Coefficient (Stock Return, Interest Rate Change)
10 years	–0.10
15 years	–0.03
20 years	–0.10
25 years	–0.14
30 years	–0.20
40 years	–0.20

We talked about equities and equity risk. We're going to move on to our next asset class, which is corporate bonds.

A corporate bond is similar to a Treasury bond. There is a higher yield. We all know what that means. In addition to interest rate risk, there's liquidity risk. Corporate bonds are pretty easy to sell if you want your money back, but there's no bid-and-ask prices in the quotes like there is with the Treasury. There's usually a little wider bid/ask spread with a corporate bond.

We'll talk a little bit about liquidity risk later on. The big risk for this bond is that the issuer might not make the payments and stiff your company. As we know, that's called a default. The way the marketplace reflects default risk is to use a higher discount rate than Treasury bonds when we determine the market value of a corporate bond. Since most of us are actuaries, we know a higher discount rate means a lower present value. In other words, corporate bonds have a higher yield to maturity when you look at the purchase price than Treasuries.

There are two pieces to these discount spreads that make up the higher yield to maturity. There's default, and there's uncertainty. The uncertainty shows up in a risk premium. With corporate bonds, we're better off, on average, than with Treasuries. Corporate bonds have a higher expected return even after I consider defaults.

A change in spread will change the market value of the bond because it changes the discount rate. Our model might also change future margins or future spreads on business. This is important if we're doing reinvestment in our model. As far as sources of information for spreads, we can get spreads from current market values by solving for them. We can discount the bond's cash flows at the Treasury rate or the swap rate, and if I add the spread to that, I'd get the current market value. Since I know the Treasury rate, and I know the current market value, we can solve for the spread. We can also get spread from published sources.

There are different methods of capturing or at least considering spreads in our models. We can assume that current spreads remain constant throughout the projection. We can grade current spreads to a long-term average. We can model the spread stochastically, although this increases the complexity of our model. It might be more than what we really want to do. Sensitivity testing could be an easier way to reveal the extent of risk. It's a common sensitivity test to decrease the spreads on the assets that we're purchasing as part of the reinvestment strategy in our model.

These are ways to model credit discount spreads. Next I'm going to talk about modeling default risk. I'm going to generalize a little bit and call it credit risk instead of default risk. There are two types of credit risk. There's the risk that the bond can default, and there's the risk that the perception in the financial marketplace of the issuer's ability to make payments might change. That's a long way of saying the bond can get downgraded. If we're more fortunate, it might get upgraded. It's common to just consider default risk in our models and ignore the risk of downgrades. Why do downgrades matter? There are three reasons.

First, there's a higher required capital with lower rated bonds. From a regulator and a rating agency point of view, we would need more surplus to be viable as a company. We can't distribute as much of our earnings to our equity owners. Second, if the bond is ever sold, we're going to sell that bond at a higher discount spread and a lower present value or lower market value. The cost of selling that bond at that higher spread needs to be considered. Third, the actual risk of default increases if the bond is downgraded. So I'm going to show an example later that demonstrates why we need to consider downgrades.

Defaults and downgrades both have a frequency and a severity. For defaults, the frequency is the incidence or how often they occur. The severity is the amount that's recoverable from a defaulted bond or the amount that's not recoverable. Somewhat surprisingly, in a lot of credit studies that are published by Moody's and S&P, the amount recoverable depends on the seniority of the debt relative to the company's other debt and obligations. It really doesn't depend too much on the bond rating. For downgrades, the incidence reflects how often defaults and downgrades occur, and the severity is how many rating classes the bond gets downgraded.

Generally, the credit risk is modeled deterministically. We deduct a few basis points of spread in our models or a piece of each asset defaults. With stochastic credit risk models, defaults can be correlated with interest rates or with economic indicators. Stochastic models can overcome some of the weaknesses of deterministic models, but they add a lot of complexity. We could also use sensitivity testing in our deterministic models as an easier alternative than stochastic modeling. The big, implicit assumption that's often used in all types of credit risk models is that asset defaults are independent. There's no correlation between bonds that we assume. The reality is that the bond market has sectors. The bonds within each sector are correlated, and all the sectors are correlated to the economy in general.

I'm going to present an example of a stochastic credit risk model that demonstrates why we should consider downgrades as well as defaults. We can take the results from a model like

I'm going to present and use it as deterministic input for our default assumptions. Table 2 is a transition matrix. My bond rating at the start of the year is in the left column. It shows the probability distribution of bond ratings at the end of the year. My end-of-year rating is along the column on the top, and my probability distribution is also shown. If my bond rating is A at the beginning of the year, there's a probability of 2.53% that it's AA at the end the year. There's an 87.8% chance that it will remain at the A rating at the end of the year.

TABLE 2
Transition Matrix

Start Year	End of Year								
	AAA	**AA**	**A**	**BBB**	**BB**	**B**	**C**	**W**	**Default**
AAA	87.37%	8.32%	0.74%	0.04%	0.03%	0.00%	0.00%	3.50%	0.00%
AA	0.86	87.30	7.72	0.44	0.08	0.06	0.01	3.50	0.01
A	0.06	2.53	87.80	5.27	0.60	0.21	0.01	3.50	0.02
BBB	0.05	0.30	6.06	83.64	5.14	0.99	0.12	3.50	0.20
BB	0.04	0.08	0.58	6.61	79.07	8.12	0.81	3.50	1.18
B	0.01	0.07	0.24	0.52	6.45	79.73	3.50	3.50	5.99
C	0.09	0.00	0.18	0.82	2.55	8.82	59.25	3.50	24.79

This table is a composite of several long-term credit studies. I used both S&P ratings, which I have on the chart, and I also considered Moody's equivalent ratings to come up with this. The W column means withdrawal. It doesn't really have anything to do with creditworthiness. It's just that the debt is retired due to maturity or retirement of the debt due to a merger or acquisition.

I'd like to point out two things for the following charts. There's an exponential growth of defaults as my bond ratings decline. The far right column, that I've labeled default, goes from zero, for a AAA bond to one basis point for a AA rated bond, to two basis points, to 20 basis points, all the way up to 599 basis points. So there's an exponential growth rate of defaults as ratings decline.

Second, bonds are more likely to be downgraded than they are to be upgraded. The numbers to the right of my diagonal sum to more than the numbers to the left of my diagonal, in general. That's not true for the whole chart but, in general, bonds are more likely to be downgraded than upgraded. In general, the numbers to the right, on my diagonal, sum to a larger number than the numbers to the left. Occasionally, I see people using this far-right column for their default assumption for cash-flow testing. That's not correct. We really need to consider the migration into the other bond categories, and we'll see that in our results.

Chart 6 shows my first point again, which is that default rates increase exponentially. We can't make out the bars for the investment grades, but for the noninvestment grades they increase very sharply. Chart 7 shows the probability of a downgrade is more likely than an upgrade. I think the downgrade is gray, and the upgrade is light gray. The point I wanted to make is that a bond is three times as likely to be downgraded than it is to be upgraded.

I'm going to do a stochastic modeling example. There's a large portfolio of 4,000 issuers. Twenty five percent of the bonds are in each investment grade rating class. So, 25% are rated AAA, 25% are AA, 25% are A, and 25% are BBB. I'm going to use the corporate transition matrix that was on the previous table. I'm going to assume that 35% of defaulted bonds are recoverable. In practice, that's an approximate average value for a group of credit studies. Our first test will be a buy-and-hold investment strategy. We're going to buy the bond. We're going to hold onto it until it either matures through that W column in my chart or at default. I'm going to replace the defaulted or matured bonds with a bond that's in the same original category. Since I have 25% of the bonds in each category, I'm just going to go back to the transition matrix. I have 25% in each category.

If I look at defaults for my four investment grade categories, AAA is zero, AA is one, A is two, and BBB is 23. I add those four numbers up, and I get 23. I have four categories, and if I divide 23 by four, I'll say that it is somewhere around six. Because I know 35% of my bonds are recovered, that means 65% default. If I take 65% of my six, that's about four. So I'd expect to lose around four basis points a year due to defaults in my initial portfolio. I have run this model.

Table 3 shows the default costs I get with a buy-and-hold strategy. The first year we get four basis points. That's a number we just figured out in our heads. The one thing we notice is that it increases every year. The reason it increases is that bonds get downgraded in the transition matrix. Two reasons it increases is because bonds are more likely to be downgraded than they are to be upgraded. As the bond defaults, and as the downgrades occur, the default rate increases exponentially.

TABLE 3
Default Costs by Year: Buy and Hold Strategy

Year	Defaults
1	0.04%
2	0.07%
3	0.09%
4	0.12%
5	0.15%
6	0.17%
7	0.19%
8	0.21%
9	0.23%
10	0.24%

One other note for this study is that I started with a large portfolio. With a small portfolio, there can be bigger ranges and larger standard deviations in a small portfolio that companies might have. By starting with a large portfolio, I can really just present averages and not ranges to you. Let's say we ask somebody in our company's investment department what the company's experience is with bond defaults. They say there have never been any defaults. We've never had a bond default. You think that's a pretty good job with a large portfolio that our company has. You guys do a better job at picking bonds than the credit rating agencies do. They say, "No, not really. When a bond is downgraded to noninvestment grade, we reevaluate that, and sometimes we end up selling the bond before it can actually go into default. That's how we've managed to have no defaults."

Does that mean that our default costs are zero now? No. It means we need to run this model with a different strategy. Our buy-and-hold strategy won't work for this. We're saying, instead of buying and holding, our investment department wants to sell some of the noninvestment grade bonds. Take another situation. Say a rating agency says your company has too much exposure to noninvestment grade bonds. In these cases, we won't get as many defaults because our bonds aren't allowed to fall into the high-risk category. The catch is that by selling the bonds in the higher risk categories, we're selling them at higher credit spreads, and higher spreads means lower prices. We lose money that way.

I'm going to do a second test that shows the cost when we sell bonds at the higher spreads. These are the spread assumptions I'm using. It is late July data. The spreads are different now. I also assume a duration of five, which is typical for an insurance portfolio. By selling noninvestment grade bonds, I'm going to use a strategy where anything that's noninvestment grade gets sold. So anything that gets downgraded to BB or a worse category gets sold. By selling these noninvestment grades, we incur losses. Chart 8 shows this in graphical form. Initially, it's much more expensive to have the strategy where noninvestment grade bonds are sold. If we look at year one, we hold everything. It's a four basis-point-cost. If we sell the noninvestment grade bonds, it costs 25 basis points. Ultimately, by the end of the 10 years, the costs are pretty close together—24 versus 28.

Buying and holding is reflected in the light bars in the graph. It's cheaper, on an expected basis, in terms of credit risk. The light bars are shorter, showing that there's lower default cost. It does have a higher risk, which I haven't presented here, and it does have a higher standard deviation for a smaller portfolio. You are more likely to actually have a bond default and incur losses. There's higher risk with a buy-and-hold strategy. By selling bonds that are below investment grade, I do have a higher average cost, but I reduce that risk. I'm less likely to have a bond default if I do that. These two strategies are pretty extreme examples. Most insurance companies fall somewhere in between the two. I've talked about corporate bonds and credit risk.

Let's move on to our next asset class, which would be a callable corporate bond. There's a credit risk with this. In the case of a corporate, there's interest rate risk, and, on top of that, the issuer can redeem the bond early. We see two time periods in the quotes. The 2004 date is the first call date. The 2009 is this date at maturity. Since the issuer can redeem early, there's prepayment risk. The issuer can pay back the money early. They'll do that to their advantage. The prepayment risk is related to interest rates. If interest rates drop, the issuer can redeem the debt early and refinance the debt at lower interest rates. So with this type of bond, there's uncertainty in the timing of the cash flows.

There are two main types of models used to model prepayment risk. There are interest rate or spread-based models, and there are market value or price-based models. Spread-based prepayment models compare current interest rates with the interest rate that's available on a new bond issue. It would call the bond if the new interest rate is sufficiently less than what the issuer is paying on the bond currently. We can adjust these interest rates to reflect whether there's a call premium on the bond or where the issuer needs to pay more than the par value to call the bond. Price-based prepayment models compare market values. They compare the market value of the bond, assuming it's never called, versus the call price or the amount the issuer has to pay to redeem the bond. If a company can buy back their debt for less than what it's worth, they'll probably take advantage of that.

The next asset class I want to talk about is CMOs. In this case, we see the issuer as the Federal Home Loan Mortgage Corporation. People call it Freddie Mac, for short. It's a government-chartered corporation. It buys mortgages, packages them together, and sells them to institutional investors like insurance companies. The gold program is a description of the types of mortgages. What that tells me is that all the mortgages in this CMO are fixed-rate mortgages. The maximum size is slightly more than $200,000. There are specific underwriting requirements that the borrowers have to meet in order to borrow money for their loans to be eligible to go into this program, and there's a special servicing program to get the money to investors. That's our company. The price, again, is in 32nd's, like a Treasury bond. To model, this we want to model prepayment behavior that reflects a group of borrowers, like the underlying mortgage holders. As with callable bonds, we can use spread-based and price-based models to model CMO prepayment. In addition to the spread-based and the market-value-based models, there are what are called

Public Securities Association [prepayment model] (PSA) and conditional prepayment rate (CPR) models that are used with CMOs. They have a different distinction of the type of model. There could be spread-based and market value-based PSA models and spread-based and market value-based CPR models. PSA is a base prepayment scale published by the Public Securities Association. Prepayments are expressed as a percentage of that scale.

A good number of practitioners think that that scale of prepayments is getting outdated. As such, CPR models are used. CPR means conditional prepayment rate, which actuaries would just call an effective prepayment rate or an effective rate. With our CPR model, our independent variable could be the difference in interest rates or the ratio of market values. One market value would assume that the mortgage is repaid, which would just be the outstanding principal on the mortgage, and the other one would be the market value, assuming that the mortgage is held to maturity. The independent variable and the dependent variable would just be the effective prepayment rate.

There are also CMO prepayment models. Beware of changes in the rules. What does that mean? It's difficult for models made with past data to correctly predict the future. Our failures in our models result from changes in the rules. As an example, mortgage liquidity greatly increased in the early 1990s. No point or no closing cost mortgages came around. That was the rule. Before that, you had to be an actuary to figure out if you were getting a better deal or not by refinancing your mortgage. When those mortgages came out, it wasn't hard anymore. In 1993, interest rates dropped quite a bit and prepayments soared. Our models that were based on the past data were very wrong. They didn't predict that level of prepayment.

What are the rule changes for the next decade? I think there's going to be an even more liquid mortgage market. I think there's going to be even easier and more convenient refinancing. We can reflect this in our models. We can't guess what the changes in the rules are going to be, but we can reflect this in our models by doing sensitivity testing.

Another asset class is commercial mortgages. Since they're not publicly traded, I don't have a quote up here. Since they are not publicly traded, they're hard to sell. I'm going to talk about liquidity risk, and I am going to talk about liquidity risk of specific assets in our portfolio. There's a current topic of insurance company liquidity or enterprise-wide liquidity risk management. I'm not going to talk about that. I'm just going to talk about liquidity risk of assets in our portfolio. There's not much liquidity to mortgages, but there is some. Some companies have sold or disinvested mortgages by packaging them together into what's called a collateralized loan obligation and selling interest in that to effectively sell their mortgages.

The easiest way that we usually capture liquidity risk in our models is to just avoid selling assets that aren't liquid, like mortgages. That works well if there's not a lot of asset sales in our modeling to begin with. We can increase the spreads to add conservatism for the nonliquid assets and try to reflect them that way. We also can model active management. For example, if I see there are too many nonliquid assets in my model, I change my investment strategy going forward to just buy liquid assets and try to put the portfolio back in balance.

To summarize, the main investment risks I've talked about are market risks, which include interest rate risk and equity risk, credit risk, prepayment risk, and liquidity risk. The main risks I have not talked about yet are currency risk. That's where foreign exchange rates can change. The values of assets that I have denominated in a foreign currency can change. If I have an asset that's in pounds, francs, pesos, or maybe Disney dollars, the value in U.S. dollars can fluctuate. We can use deterministic and stochastic modeling to model currency risk. Volatility risk has to do with both interest rates and the equity market. If we have equity-indexed products, it's very important to consider volatility risk.

I'm not going to talk too much about that because I think there's another session on equity-indexed products. Most of the people interested in that aren't here. The volatility is important if the company uses any type of interest rate or equity hedging involving derivatives. The price of the derivatives increases as the volatility increases. The price of the derivatives or the price of the assets I'm using to hedge increases as the volatility increases. The cost of my hedging strategy depends on volatility. We can use sensitivity testing, or we can use stochastic modeling to test volatility risk.

A good example of political or legal risk is the utilities in California. Pacific Gas and Electric was an investment grade bond. In California, electricity was deregulated or at least partially deregulated. Now the bondholders are fighting in court to get back part of their investment that they made into those investment grade bonds. We usually fold political risk into credit risk. Event risk considers corporate restructuring. An example would be the leveraged buyout of RJR Nabisco where a lot of investment grade debt suddenly became noninvestment grade. We fold investment risk into credit risk also and consider it there in our modeling.

I've talked all about risks. I'm going to talk briefly about some other modeling considerations—calculating market values, combining assets and liabilities, data, and assumptions. There are methods of calculating future market values within our model. Most of us project the asset cash flows, discount them at the risk-free rate plus the risk premium. That compensates companies for explicit risk, like default and liquidity, and for the fact that defaults are not known with certainty. This type of valuation is most common for assets that have fairly certain cash flows. If the cash flows aren't certain, we can project the cash flows under different scenarios and use the risk-neutral models that we talked about to value these.

There are a couple of less-often-used valuation methods, which would be used to modify the future cash flows to reflect risk and then discount at the risk-free rate. We can also modify the future probabilities to adjust for risk and discount at the risk-free rate. That can give us the same answer as this first method of using the risk-free rate plus the risk premium. If we do have assets with uncertain cash flows, a common approach is to use option-pricing models. There's Monte Carlo simulation. There's also quasi-Monte Carlo simulation, which a lot of our software might use, which really takes some representative future scenarios to value the asset. There's closed-form option pricing. Oftentimes, we can value our assets and value options using a formula. There's also lattice form option pricing, like a binomial lattice method, that's used to value those assets.

Far as combining assets and liabilities in our projection, a lot of us use an integrated software package that does this. Our asset projection software and our liability projection software are different software. Assets are projected into a different system. Things we need to consider here are asset sales. What if our model needs cash to pay liability obligations? We can bring the market value of the asset into the model that projects our liability so the asset can be sold. Another simple way is to sell part of the entire asset portfolio, or we can borrow money to raise cash instead of modeling sold assets. If we do that, we need to do that carefully in our model so we don't have unreasonable leverage or arbitrage.

New asset purchases have to be modeled as part of our investment strategy, as well as the interaction with our liabilities. An example would be, if I had portfolio interest rate crediting, the liabilities affect the assets, the assets affect the liabilities, and we get a little bit of a circular problem that I need to resolve. There's also software and data transfer issues. Getting one piece of software to accept data from another can sometimes be a challenge for us.

For sources of data and assumptions, we internally have our own investment accounting systems. That's where we get most of our information for asset modeling. We can also get some information and assumptions from cash-flow testing regulations, New York Regulation 126, and the Actuarial Opinion and Memorandum Regulation. There's some very broad, general guidance in the Actuarial Standards of Practice 7 and 14. We can ask the brokers who sell the securities to our company. If the brokers do a lot of business with your company, you've got some leverage asking them for information. They're very helpful in providing that. We can also ask the issuers of the bonds or the mortgage agency that sells CMOs. They provide websites that have information about the investments to help out investors. There are software vendors that make our asset/liability modeling software. There are also outside sources, for example, Bloomberg or FRED, which is the Federal Reserve Economic Database. That concludes my part of the presentation. I'd like to open up the session for questions.

MR. STEPHEN N. STEINIG: You spoke briefly about modeling the spreads. Could you say a little bit more about how you go about doing that in practice. In particular, can you address the case of how, when you are working on a 10- or 20-year model, you would keep the spread assumption and the default assumption consistent with each other?

MR. HOUGHTON: For modeling spreads, a pretty common assumption is to assume that the spreads that exist today stay in place or stay constant through the projection period. If we use a default assumption that's constant, then at least we have consistency between the two that way. Another approach is to grade the spreads towards a long-term average. If we think today's spreads are an anomaly, we want to make our best guess that they'll change over time. We can grade them down or up in our model to reflect that. Reinvestment in a model is more conservative to grade them downward. It might be a little bit of an aggressive assumption to grade them up because we think they're going to increase. One approach that companies can use is to do a sensitivity test, or just run another set of projections using lower spreads. We can determine how those look compared to our baseline projections. There are stochastic models for spreads. They can be derived randomly and still be used with a stochastic default model that's consistent with the spread model. That is more than what a lot of us would want to do, but it is possible. Any other questions?

MR. HAMBRO: Scott, I have one question about spreads. We're in an environment right now in which the Treasury rates are very low, and as the Treasury rates have declined during the year, the absolute yield on corporate bonds has not declined that way. Actually, it might have increased a bit. What are companies going to need to do for cash-flow testing and other ALM projections to reflect that?

MR. HOUGHTON: We have interest rate scenarios that change going forward. A conservative assumption is to perhaps grade down the spreads in our reinvestment strategy as a sensitivity test. It's really up to what the appointed actuary feels comfortable with in that projection. I said I'd feel more comfortable grading the spreads down as interest rates increase.

MR. HAMBRO: If there's a mean reversion assumption in the actual stochastic interest rate generator, then rates are going to increase. They will probably keep the spreads at the same level, and it might be appropriate to adjust the spreads accordingly.

CHART 1
Treasury and Swap Curves

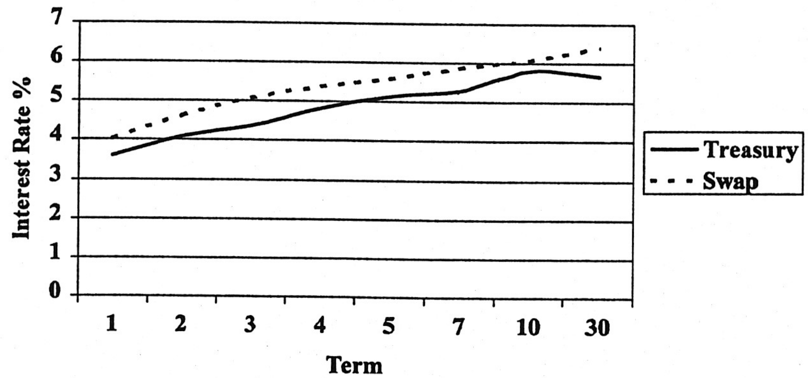

CHART 2
Stochastic vs Deterministic
Test NY 7 Results

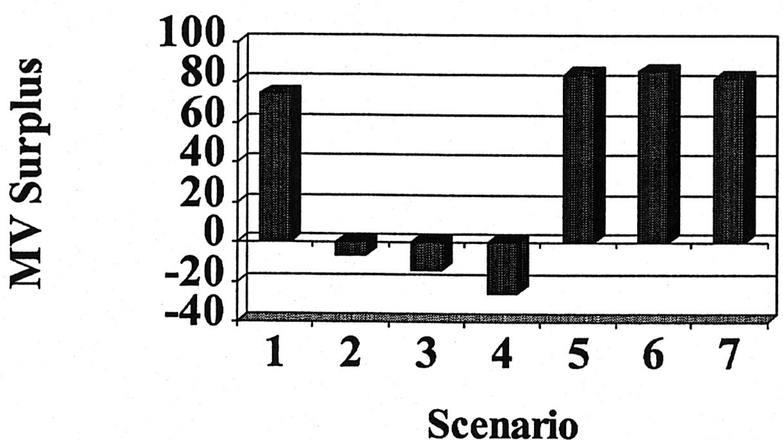

CHART 3
Stochastic vs. Deterministic
1000 Stochastic Results

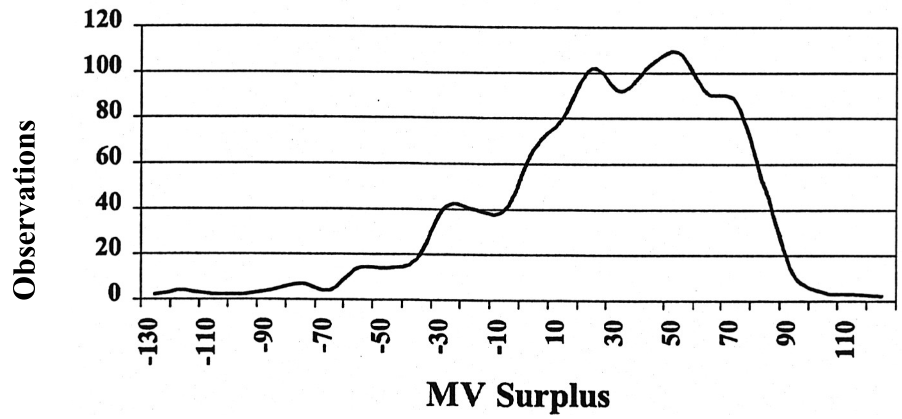

CHART 4
Stochastic vs. Deterministic
Combined Results

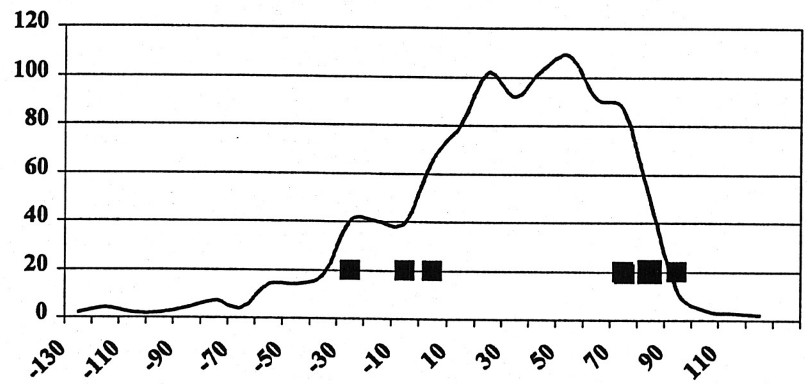

CHART 5
Distribution of Stock Market Returns
1939–2001

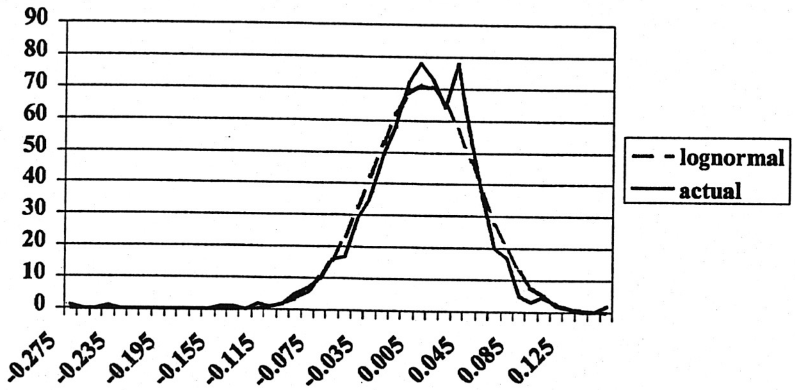

CHART 6
Transition Matrix Default Probabilities

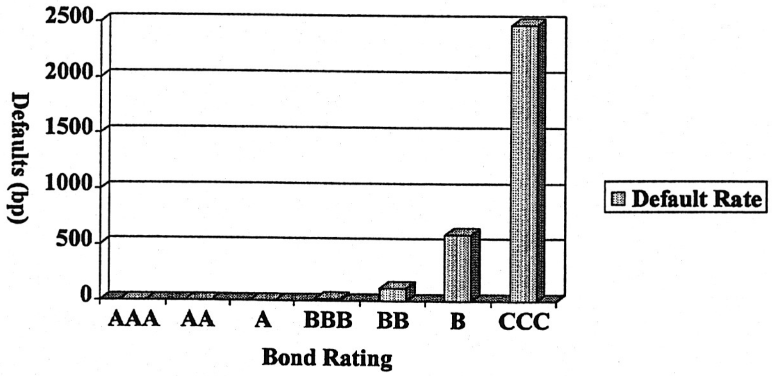

CHART 7
Transition Matrix Summary Probabilities

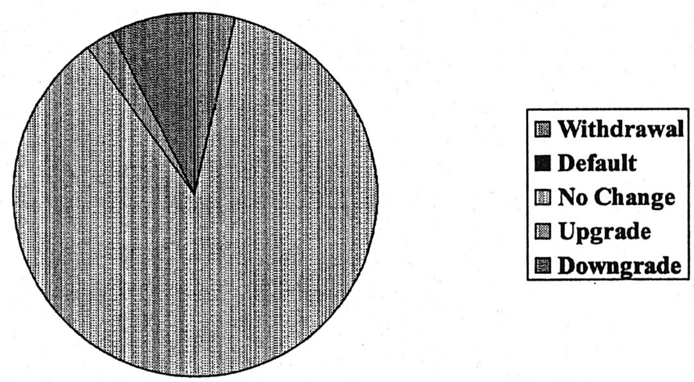

CHART 8
Default Costs by Year