2004 Valuation Actuary Symposium*

Boston, MA
September 20–21, 2004

Session 5TS
Asset Modeling Concepts

Moderator: Ellen Cooper

Panelists: Don P. Wilson
Ellen Cooper

Summary: The need to model assets in insurance companies has gone well beyond compliance with cash-flow testing requirements. Robust asset modeling techniques are needed in order to quantify risk exposures appropriately and to comply with a widening array of liability valuation and capital adequacy calculations that rely on stochastically modeling the interaction of liabilities and the assets that support them.

MS. ELLEN COOPER: We're going to spend about 90 minutes today talking about asset modeling concepts. I'm going to be your moderator, and I'm also going to be a speaker, which means I'm going to be introducing myself. It's always intimidating when you see some of your current clients in the audience and especially when you're working with them on some of the related topics. In any event, I want to introduce our first speaker to you. He probably will be a new face to many of you. This is Don Wilson, and Don will be speaking on a bunch of technical issues obviously around asset modeling. He comes to us from the United Kingdom, and he has 15 years of experience in a whole host of things.

He works for Deloitte in the Hartford office and prior to that was based in the Deloitte London office. He has an extensive background in asset/liability modeling (ALM) and stochastic modeling. He does a host of client work, as well as a significant amount of R&D. If you want more information about his background, he was good enough to give me a page or two, and I'm going to spare the details for now, but if you're interested, you can take note of his e-mail address, which is donpwilson@deloitte.com.

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

MR. DON P. WILSON: I'd like to start with reading a slide to you. It's not in your handouts. It's a new one, and I'm not going to explain where I got it. Most people see risks and try to avoid them. The actuarial profession was created by people who want to manage the financial consequences of death, a risk that cannot be avoided. The profession has moved on from its origins more than 150 years ago. Actuaries now work in areas where the management of assets and liabilities, especially over the long term, is assisted by a combination of mathematics, economics and finance. Despite the popular myth, actuaries cannot predict the future or make unshakable forecasts. We have never been able to. Actuaries build financial models, some as familiar as a spreadsheet projection, others based on stochastic random simulation processes, all founded upon assumptions about uncertain future events and all designed to throw light on alternative future circumstances so people can plan for a range of outcomes.

I thought that was pretty pertinent, not just to this session, but probably to all sessions and what we do. Actuaries tend to be much stronger by training in the modeling of liabilities than they are in assets, and one thing we're going to do today is to try to move you forward on the asset side.

I'm going to share with you some thoughts on what makes a good stochastic generator, and then go onto something that's been a hot topic recently. I'm sure you're now pretty competent on the difference between real-world scenarios and risk-neutral scenarios, and for this I'm going to borrow some slides from a presentation that was given at the last SOA spring meeting. Then I'm going to talk a little about deflators. And finally, because it was on the program, I'm going to add a few words about validation of assets.

The first thing I want to talk about is what makes a good stochastic generator. There are two areas of that. It needs to be realistic, and I'm going to cover what I mean by realism in some greater detail. It needs to be economically coherent. Remember from the introductory slide, we're not just mathematicians now. We have to understand finance; we have to understand economics; and our models have to cover all three of those things. It has to be statistically sound, otherwise it won't be credible, and it has to be scalable. For instance, in scalability, if you're modeling variable annuities (VAs), you need to be able to model a number of different equity funds or investment funds. So you need to be able to cover the variability around that.

I'm not going to talk much about scalability. We'll concentrate on the realism aspects. On the realism aspects there are really four fundamental economic principles, and I'm going to define all of these principles or attempt to define them at least in terms that I understand and explain why they are important: efficient markets, arbitrage-free, full-term structures and equilibrium pricing.

First I'll discuss efficient markets. Chart 1 is a graph of the Standard & Poor's (S&P)

500 Index up to September 1999. Wouldn't it be nice to have known what was going to happen next? But the thing about an efficient market is that all the information is already in the price, and there is no way that you can tell what's going to happen next. What happened next: The market took a considerable dive and has risen a bit from there, but is still pretty flat.

With efficient markets, all publicly available information is already in the price, and it's what causes the market to be volatile. We can see that the market's volatile by the way the prices go up and down in a pretty erratic fashion. If all the information was not available, and it wasn't an efficient market, we'd see the prices moving a lot more smoothly. So, it's important to have this concept of efficient markets because if you don't have that concept, then there's a risk that when you are seeking to optimize your investment strategy, you will find solutions that are not actually available to you in practice. In that case, you may be led down the path expecting to make higher gains than are really realizable. Also, from that point of view, you might end up in hedging strategies, which are not going to work in practice. So, I think the key thing here is that even if a market isn't 100 percent efficient, it is important that the stochastic generator that you use is. You must avoid the risk of picking a scenario that is unrealistic.

The second concept is arbitrage-free. If we take the last four months and if there was an arbitrage possibility in currencies, then we would expect perhaps at some stage to be able to take our dollars, buy yen, sell yen and buy pounds. So if we look at the pound-to-dollar scale, there is a perfect market and there are no arbitrage opportunities in currency. Obviously, arbitrage opportunities might arise in other stocks that are not so widely traded, but it is important for a model that you do not assume arbitrage because, again, you could end up with investment strategies that are just not realizable in practice.

When I stared at Chart 2 for a moment this morning, I thought this is like one of those optical illusions. You stare at it for long enough, and it suddenly flips to look like something else, but it's actually just a chart of the progression of a yield curve over one simulation coming from a stochastic generator. It's generally an increasing yield curve with outstanding duration, but it wanders all over the place, and it seems in this simulation to have a bit of a periodicity about it. Interest rates go up, and then they come down, and they go up, and then they come down, which is what you'd expect. If you price them incorrectly, you may make the wrong pricing decisions. You may make the wrong reserving decisions. There's a whole series of things that may go wrong in the model. And, indeed, you can also be left with arbitrage possibilities. So it's another thing on the list of at least ideal requirements for a stochastic generator that you do include a full-term structure in it.

The fourth item is equilibrium pricing; by this I mean that the price will adjust to the point where in any market your supply equals demand. That's basically how pricing is set in an efficient market. So, in the context of stochastic generators and in the markets that they are modeling, we're looking for the level of risk premium at which you get a balance between supply and demand. If you want to earn more than a risk-free rate, then you need to take a risk, and the more risks you take, the higher the expected return would be. You get an upwardly sloping yield curve. This is important because you need to be consistent in your generator with what the market is saying in terms of the relative riskiness of the different asset classes, and you need to be objective with your levels of risk premiums. If you choose some arbitrary levels of risk premiums, you risk them being inconsistent with other aspects of the model, and, therefore, you risk your modeling coming up with strategies, again, which are overly optimistic.

So these are all principles that are important if you want to be able to be confident in the results that are coming out of your model. The model also needs to be statistically sound. The stock markets and interest rates do not move smoothly. They are fat-tailed. They have skewed distributions, and you need to allow for that in the modeling of the stochastic outcomes. The stock markets also jump at times, and ideally you would include jumps from that point of view. These are complicated things to include in generators, but it's where the world needs to go if we're to have confidence in the results. To make sure that the inputs at the starting point for a generator are consistent, particularly the risk premiums, initial yield curve, the volatilities, everything needs to be consistent at the start.

Those are just a few thoughts on stochastic generators. I'll set that in context in the next section where I'm going to cover some aspects around real-world and risk-neutral scenarios. We want for some purposes to produce simulations that reflect what we think might actually happen in practice, but to set those sorts of assumptions can be very difficult. What level of expected returns might we expect on equities in the future—about 8 percent, 9 percent or 10 percent? Even 10 percent starts sounding a bit high, but the simulations produced by the Academy of Actuaries for use with the risk-based capital (RBC), C-3 Phase II requirements have an expected return of 12 percent or more, and that's because those scenarios were based purely on a historical study. The mean of history over the last 30 years has been 12 percent or more returns. But should we assume that we're going to go back to some periods like in the past and assume a lower level? These are very difficult decisions to make and make a tremendous difference on the results.

So, you get some very different results based on the different style of model stochastic generator you use, and even if you're putting the same means of standard deviations because for the stochastic generator what you're most interested in are the tail ends of the distributions. Those tails can be very different. So, to give some examples on that, the pricing actuary might well take a more optimistic view than the valuation actuary. It might assume a 10 percent return, a volatility of 16 percent and use a simple stochastic generator. The valuation actuary takes a more conservative view. I guess you're all conservatives on that basis here today, this being a valuation actuary symposium, maybe a slightly different model.

Certainly, if we look at what the Academy of Actuaries has produced for VAs, it isn't particularly concerned with the return volatility. They're just concerned with how you fit the tail of the distribution to make sure that you have enough extreme negative events so that your reserves are suitably prudent, but all three views are supposedly real world, and they all give an impression of plausibility. They look as if they could be right. They can all be expressed in terms that you can understand. The key thing about these real-world scenarios is it means you can draw conclusions directly from them.

You can look at the outcomes and say, well, 5 percent of those go negative by so much. That actually means 5 percent probability that that occurs and that gives you the ability to set reserves based on percentiles or conditional tail expectations, things like that, but there's nothing actually there to help you determine which one is correct. There is nothing around to say that anything is going to be correct. We cannot properly foretell the future. And that's why the initial principles that I went through earlier are important, because we can't tell directly from the results of the simulations whether this is sensible or not. We have to revert to the underlying principles in the stochastic generator and be sure that they are sensible and coherent in order to be able to believe the results of the stochastic work that we do.

Let's now switch to risk-neutral simulations. I'll give a very brief definition of them. The thing about risk-neutral funds is all the investments produce the same expected return. There is no risk premium in that.

Equities: We'll take a simple binomial model. The price may move from 100 either down to 95 or up to 125, and the question is what is the probability of movement, and how do we value the equities? The real-world view will say, well, it's a 50 percent likelihood of going up and 50 percent likelihood of going down. And that's consistent with an equity risk premium of 5 percent, if you do the math, but the question is how you then value that equity asset with those two probabilities. We have to find a discount rate to value the positive outcome. When you get 25 out of 100, and that's a common problem that I'm used to from the U.K. of embedded values, the choice of the discount rate is the biggest, most difficult issue of all. Without that discount rate, it's very difficult to put a value on it.

So the risk-neutral solution to that is to solve for probabilities, to bring the equity risk premium down to zero, and if you do the math on that, you find that the outcome of 125 is a probability of one-third, and the outcome of 95 is a probability of two-thirds. So, we can then value the outcome by taking one-third on the 25 and discounting at the risk-free rate, and that gives a value on the equity assets. That's how the asset markets are valued, and that's the basic risk-neutral approach to valuation. The key thing here is all assets are discounted at the risk-free rate. There's no subjectivity of choosing a discount rate. We get a good measure of prices, and these are market consistent prices. In other words, if we're putting a value on a liability, it will be consistent with the way values are put on the assets. But this model, because it's generated, has used some artificial probabilities in order to come up with a valuation. It cannot be used for the purpose of looking at probabilities of outcomes because the probabilities are distorted in order to give you the valuation. So, the real probability's 50/50, but for valuation purposes it's one-third or two-thirds.

I'm going to very quickly pass through a case study. So, let's just look at the values on it. This is a guaranteed minimum income benefit (GMIB) and a guaranteed minimum death benefit (GMDB) in a VA, and we've just set up a little generator to do some simulations on the outcomes of this model. If we look at the distribution of outcomes, we get a typical chart. The death claims are a small number of simulations, which have a very high payout, and a few that have no payout at all. If we take the mean value of those—and these are real-world simulations—then we get to one level just over 50. The risk-neutral approach puts a higher value on it because it has weighting. It weights more heavily to the extreme scenarios.

It can look at the same thing for the income benefit where the distribution is more extreme. A large proportion of the outcomes have no payout. A few have a very high payout, and again, the risk-neutral value will put a much higher value on than just taking the mean of the real-world scenarios. You could get around this with the real-world scenarios by taking a conditional tail expectation, so you take the average of the last 10 percent. But the question that you're still faced with is: Is it the last 10, the highest 10 percent or 20 percent? How do you determine which ones form a reasonable value? So, using the risk-neutral approach to valuation does provide a mechanism for producing reasonable values on the various benefits.

So, what this is showing is that using risk-neutral simulations is very different from using real-world simulations, but the risk-neutral ones are tipped to put a greater weight on the negative outcomes. Markets are generally risk averse, and, therefore, that's what is demonstrating that feature. But all of these aspects are still quite difficult to resolve in terms of choice of volatility, and the number one variable is policyholder behavior. If I could just digress on policyholder behavior, I've been working a bit on a model looking at the different policyholder behaviors and looking for VA contracts and looking at the capital requirements. I took two extreme behaviors, one a passive behavior.

Someone has a guaranteed minimum withdrawal benefit (GMWB) but doesn't do anything with it. The other manages it to try to optimize the return. I found with the passive behavior there was no additional asset requirement in RBC, C-3 Phase II, but with the active one the sky was the limit. If you put in active policyholder behavior and astute policyholder behavior, you found that the capital requirements were 5 or 6 percent of the premium. It's a tremendous range of possible outcomes, and so investigating that and coming up with sensible algorithms for that is going to be very important. That was a slight digression.

So, I've covered a little bit about real world, risk neutral, and the current thinking is you have to make sure that you use the right type of generator and the right types of simulations for different purposes. In principle, that would mean that you'd need two stochastic generators and two sets of simulations in many of the projection models, but there is a solution that avoids that—deflators. If you could use a stochastic generator that produces deflators (path-dependent stochastic discount rates), they are weighted toward the poorer performing scenarios. So, in a sense, from that point of view, they have the risk-neutral characteristic, but you use the real-world scenarios, and then you use these deflators to discount with, and they will then give you a risk-neutral type valuation. So, if you can use deflators, then you can use the same scenarios for the projection part of the modeling as you do for the valuation part. We have a 50/50 possibility of going up and down and a cost of guarantee, simple model. In a risk-neutral world, to get to those probabilities of up and down for risk neutral, we have to solve for two conditions.

First, if it's a binomial model, the probability that the price goes up plus the probability that it goes down must equal one. It has to go one or the other way. If you take the initial value, the probability goes up times the amount of the increase. Plus, the probability goes down and the amount it goes down, you should get to the expected return of the portfolio increasing at the risk-free rate. If you solve those two equations, you can calculate the probabilities that you need, and then if you apply those probabilities, you come up with a value or a cost of the guarantees, in this case 3.17. So, that's the risk-neutral approach to it.

Basically, the use of risk-neutral probabilities discounting at risk-free rates is exactly the same as using real-world probabilities and discounting with deflators, and what we've done is moved the riskiness from the probabilities and the projection into the discount rate. Of course, the question is: How do you get to these deflators? They are an inherent part of some stochastic generators. They come from within the center of it. So, they have to be a property of the stochastic generator, as well as producing interest rates, equity returns and inflation returns—all the outputs. They also need to produce these deflators as outputs, and if they can do that, then you can apply them as your discount rate rather than the risk-free rate in a risk-neutral world, and that will give you the market-consistent values.

You can read more of that in the July issue of *Risk and Rewards*, the SOA Investment Section's newsletter, in which I wrote an article that explains deflators in much greater detail. So, you can read up about that in your free time if you want to know more about deflators. If you have the deflator probabilities, then you can calculate the cost of guarantee, and you'd expect it would come out to the same amount as the risk-neutral approach, but this time we're using real-world probabilities and calculation.

Now, I'll come to the last part of my session here this morning. I have a few very simple words on validation, sort of big jump in topic here. This is really speaking from my practical experience, and the first thing that I was very careful to do is to make a rigorous split between the stochastic generator, which is a model of the economy, and the business projection or projection model that you're working with. There's no insurance company that we would assume is dominant enough to actually affect the economy by its actions.

So we can assume the economy is independent from the business model, and it is very helpful to keep that split because it means you can sign off on the simulations from the stochastic generator before you run your business model or your projection model. That makes a great deal of sense. I had a case the other week. I forgot to do this, and it came home to me, and then I ran it, and I found the volatilities. When I did an analysis of the means and volatilities of my stochastic generator, I could see that there was a clear error in that. The volatilities were completely wrong, and I'd go back and rerun it. So you really should sign off on the stochastic generator to make sure that you're comfortable with those assumptions before you feed them into the second stage in the model.

In the projection model, if you're running it stochastically, it's very important to check individual simulations. Pick some good ones. I tend to seek out the worst performing scenario and check through the detail of the calculations very carefully to make sure it's working properly. You can check everything in a stochastic projection, but there are certain things that are more important to check, and if you look at the way the calculations work out for some of the worst simulations, you can make sure that you debug the model. It's also important in a stochastic world to only change one thing at a time. If you change too many things and then rerun the model, it's very difficult to interpret what came out of it. You really get stuck as to whether this was caused more by one effect or the other effect. You need to change one thing at a time. It's easier to interpret and easier to learn. We're, all of us, learning our way into this stochastic world, and it's much easier to learn more if you just move one thing at a time, work out effect, and each time your understanding goes up and your confidence goes up.

Also, it's very important to document what you're doing before you send the results up to your manager, ideally before you even run the model. Just writing things out helps avoid silly mistakes. There'll always be compromise in the stochastic world. These are only models that we're developing. You can't do things exactly. Things take too long to run. We have to be careful to avoid over-interpreting the results from the model. The model is only as good as the assumptions put into it or rubbish in-rubbish out.

So, you start off with a question. In order to address the question you have to come up with some methodology for it, some assumptions and some data. You push it through a calculation engine. You get some results out. But, crucially, the interpretation of those results is in the context of the methodology and the assumptions that you've used. Hopefully, you'll be able to decide something and actually do something as a result of running the model, and that might take you back to asking some further questions. So you keep on going around. Well, thank you for listening to me this morning. I'm now going to pass you over to Ellen for her part.

MS. COOPER: I am a senior manager with Ernst & Young. Like Don, I've also relocated all the way from Philadelphia to New York, where I now work in Times Square. I am also a Fellow of the Society of Actuaries. I have just passed chartered financial analyst (CFA) Level 3. So I filled out that charter application, and assuming that I don't have any traffic violations or any ethical issues, hopefully I will be admitted to be a CFA charter holder relatively soon, but I can't yet say that I am. I also have about 15 years of experience, like Don, and my background is quite similar to his. I do a lot of ALM and stochastic modeling. In addition, currently we're doing a lot of derivative-based hedging work, C-3 Phase II and the hedging implications within C-3 Phase II, economic capital and so forth. Today we're going to talk about asset modeling.

First of all we're going to talk about asset-related assumptions. We're going to then walk into market value of assets and go through some sort of basic theory around embedded options, which probably a lot of you already know, but it's always good to review. We'll then walk through duration and convexity. And then the presentation will culminate with an example of dynamic rebalancing where we use duration and convexity to do the dynamic rebalancing.

So, why do we care? Why is asset modeling important? What do we care about? I tend to see a lot of clients fairly frequently, and one of the big questions is always understanding what the implications are of my reinvestment and disinvestments strategy. If I reinvest in Strategy A, what happens to my earnings? What happens to my investment income versus Strategy B? We have to do cash flow testing. We're forced to do asset modeling. But, in addition, there's a lot of stochastic modeling that's done for risk-management work to understand tail scenarios, to understand what we can expect. Recently there's a lot of work that's been done around investment income at risk and understanding distribution. Investment income is such a key portion of operating earnings and understanding what my potential distribution is and that I could potentially miss my earnings target, not from too many withdrawals from my VAs but, in fact, from getting my investment income wrong. So controlling earnings volatility is obviously an important issue.

In terms of asset modeling assumptions, for those of you who work with asset models, you're going to be familiar with every single one of these. There are several key things that we need to worry about when we're setting up our models to try to essentially replicate what's going to happen in the future. We need to worry about our prepayments. And we need to, by the way, know which assets have prepayments and which ones don't. That's important for a bunch of reasons, and I'll talk about some of that later. But if we have, for example, a portfolio of asset-backed securities that typically don't have prepayments, we may not want to model them using Intex and going through a full Monte Carlo method where we're doing a thousand scenarios at every node and wasting a lot of runtime.

We might want to look at those cash flows and say, you know what? This really isn't interest sensitive. I'm going to replicate the cash flows by using a straight non-callable bond instead. So, we want to make sure that we understand which are prepayment-sensitive, and then we want to make sure that we understand what our prepayment algorithm is. I know a lot of companies rely on Adco, and none of us really understand, I don't think, what it does. It's somewhat of a black box, and hopefully that's fair to say. You want to make sure that you're studying how your prepayments are changing relative to your interest rates. Look at some individual seriatim assets and make sure that you're comfortable with what they are doing, and, if not, use something else because I'm going to show you later how key prepayment assumptions can be to your model, particularly if you're looking at duration.

Credit spreads also are very difficult. If you're running a 20-year projection, credit spreads don't seem to have a real relationship to level of interest rate, to even necessarily level of default risk or credit risk as we saw during the last down cycle of credit. So, it's very difficult. Do we want to use some kind of stochastic credit spreads? Do we want a credit spread that is just a straight assumption? Do we want to do something that's a multiplicative off of interest rates? And one of the reasons why it's so difficult is because for my work there's not really a clear functional form. So, I think the best thing to do is to be aware of it and to be aware of how varying credit spread assumptions will impact your models again both in terms of cash flow, in terms of market value and obviously in terms of duration and convexity.

Default risk and associated recovery rates is another place where we tend to go light in terms of our work. Again, here is some of my experience. I work for a company that audits, so I have the luxury of seeing a lot of cash-flow testing models in an audit context and working with a lot of companies in terms of asset modeling assumptions. We tend to rely on C-1 factors, or we come up with something that seems reasonable, or we use the Moody studies. It's important to try to calibrate the Moody studies to our actual experience, and one of the things that we insisted during the down cycle was to have some kind of a grading of long-term averages.

It's probably not appropriate to use a long-term average default rate when you're in the middle of a down credit cycle, and you know that you're going to be experiencing significantly more default risk for the coming years. You want to trend that over time. And then also you want to make sure that you are in contact with your investment department. If your investment department has some kind of credit metrics or some kind of Merton-model-like model, you want to make sure that there's some relationship or connection, that you're not off in a vacuum setting assumptions. Work together to set those assumptions.

What are recovery rates? They are the loss given default. If my issuer defaults, how much am I going to recover? And, again, this has been very volatile. I'm going to show you some graphs a little bit later, but we tend to make very quick assumptions, 50 percent. We need to understand our issuers, and we want to work again with our investment departments that are studying these issuer by issuer and, say, telecom industry, subordinated debt. What should the assumption be? Have different assumptions by the type of debt that it actually is and by the issuer, by the industry. It will impact our models.

Reinvestment/disinvestment. We want to make sure that we don't just throw in 10-year AA bonds if that's not really reflective of what we're doing. And on disinvestments, borrowing can be dangerous. First of all, I don't think any insurance company actually borrows when they're in a negative cash-flow situation. But you want to make sure that if you have a borrowing assumption in there that you don't have a leveraging effect in your model, that you're not borrowing at some low rate, using the negative cash flow to reinvest at some higher rate, and then you're showing all these earnings that are never going to happen because that's not really going to happen in real life.

What do we need to think about? We obviously need to think about interest rate risk. I'm going to show you a graph later, which everybody knows, of just how the yield curve moves around from period to period. The shape of the yield curve changes. And we need to understand how our assets change relative to interest rate risk and also the prepayment behavior, which part of what changes it is the level of interest rate, how that changes. Equity risk. Today I'm predominantly talking about fixed income, but equity risk obviously is going to have a major impact on any assets that have an equity component. In addition, if you're looking at GAAP earnings projections, you need to look at the various accounting policies and methods.

If you have securitized assets that are affected by Financial Accounting Standard (FAS) 91, you want to be able to look at and show to senior management alternative scenarios under alternative FAS 91. As opposed to looking at cash flow, we're looking at something statutory. You want to make sure that you're giving them reflective, appropriate GAAP earnings, appropriate investment income so that they have some indication. Finally, there is the way that assets correlate and the way that we can diversify, and that's both by asset class and also by issuer. I think we saw during the down cycle of credit that there were many insurance companies that had too many airlines, too many telecom-type issuers. So we want to make sure that we're paying attention to that as well, and if we are not well diversified, we either want to think about whether we want a program to further diversify or make sure that we have appropriate risk reflected in our models.

Chart 3 looks like my seven-year-old took a crayon and drew this. Actually, it's better than it used to look, but you can see how the term structure floats around quite a bit from period to period. Lately, there was a period of time when we were all concerned with how long we were going to go down and stay down, and how we were going to survive during the down cycle. Now there are more questions around interest rates rising. How quickly are they going to rise? What's the volatility around the rise going to be? And how is this going to impact my portfolio and ultimately my earnings?

Chart 4 shows average credit spreads. Again, this is just to give you some example, by average credit quality, how things kind of flip around. If you do a quick-and-dirty analysis and you look at credit spreads versus level of interest rate, you'll see that it's very hard to make any kind of statement or come up with a functional form to relate interest rates to credit spread. This is not completely updated because I, like Don, sort of found out last week that the session was, in fact, on. So I didn't have a chance to update this graph, but it portrays something important, which is if you look at 2001, and you look at what our default rates look like, we were in a situation of record default, and we have to be careful.

Basically, the experts are saying that we're in the peak of the down trend right now, that we're in the peak of the low credit cycle, which means that we're experiencing good credit risk, not bad. It's likely to happen again, unfortunately, and so we want to be careful to make sure that we have the right tools in place to ensure that we can do a better job of managing through the next downturn of the credit cycle.

Chart 5 is looking at recoveries and understanding the relationship, that there is a high correlation between the high level of default and the low level of recovery rates. So, those tail scenarios become exacerbated because of the fact that we have high defaults and weak recoveries in the same cycles. Chart 6 is just a picture to show you the movement of the one-year versus the 10-year note and how the slope of the yield curve has really widened over the course. Obviously, this has come back up and is not fully updated.

So, to summarize what we're saying in terms of a modeling context, we need to ensure that we have good methodology around scenario generation because what happens to our key economic variables or interest rates or equity, foreign currency, if it's applicable to your organization, is going to be key. Credit risk as well is key, and in my opinion we need to give more consideration to how we're going to project default risk and loss given default in the future. Don spoke to us about the issues around scenario generation.

Now we're going to move into the world of embedded options. We have a good scenario generator, and we have all these really great assumptions. We have good prepayment forms and so forth. So what is an embedded option? An embedded option is basically the right to take an action against another party that is embedded inside of our issue. So we're buying something or we're selling something that gives somebody else a right but not an obligation to take a certain action.

In the case of a callable bond, I'm going to talk about two examples today—a callable bond and a securitized asset. With a callable bond, the issuer has the right to call the issue. In a securitized asset, there are borrowers in a pool, and they can prepay an amount in excess of the scheduled principal payments. And then, in addition, we have on the other side that the bondholders, or the person who actually owns the issue, can convert if they have a convertible bond. They can put an issue, which is like essentially surrendering a policy, for example, if market interest rates rise above a coupon rate. And the most important thing is that option exercise depends on the level of the prevailing interest rate relative to the issue rate or the borrowing rate for securitized assets. Scenario generation is so critical to this because if you have bad scenarios, you're going to have bad embedded option analysis.

I'm going to talk briefly about the simplified method of yield to first call or yield, which essentially builds to yield to worst versus the lattice method. So yield to first call, what is it? Basically, you take the assumed call price and the assumed call date, and you look at and solve for it using a yield to maturity method. What is the yield that I need to have such that the present value (PV) of my expected cash flows is equal to my market price plus accrued interest? If I have 5 percent 10-year bond with a maturity value of 100 selling for 105.5, what do I do? I know that my first call date is four years from now, and I have a call price of 102. I have cash flows. I have eight coupon payments of $2.50 each for six months, and I have a final call value of 102 in eight six-month periods from now.

I basically set that up, and I have a one plus y to call on the bottom. That's my unknown, and I solve for that yield. If I have a yield to worst method, which is typically in most of our asset models, what I do is I look at this under every possible call date, and I come up with the lowest yield, the yield to worst, and whatever the price is at that point, that's my yield to worst. It's basically a single-path method. What's wrong with this? Well, what's good about it, first of all? It runs fast. We like it because it runs fast, and it's simple. But what's wrong with it is that we assume, first of all, that all cash flows are reinvested at this yield to call rate, which is not what's going to happen. We assume that the issue is going to be called on the assumed call date. Now, for yield to worst, obviously, we're looking at the worst possible. And if we have a bond that has a yield to call method versus a yield to maturity method where there is no call, they're not directly comparable.

Now, to contrast that we're going to just talk for a minute about option-adjusted spread (OAS), and then we'll talk about lattice method. So how do we calculate a spread? The simplest way that we calculate a spread is we look at the yield to maturity of a corporate bond, and we compare it to the same point on the Treasury. We take the difference, and we say the nominal spread. The spread is 125 basis points. In contrast to that, we can go a step further, and we look at what we call the z spread, or the zero volatility spread. What we do is we come up with the spot rates. Once we have the spot rates, we look at the cash flows that line up with each of those spot rates, and we solve for a static spread such that the market value equals the cash flows divided by the spot rates plus the z spread. So we take it a step further.

FROM THE FLOOR: What does off one point mean?

MS. COOPER: Off one point means that essentially the yield to maturity method solves for a single point in time as opposed to the spot rates where you actually have different spots at different points in time. Yield to maturity takes one yield and discounts all cash flows by a single yield. We take the z spread one step further, and we go to an OAS. The difference between z spread and OAS is that now for OAS we know that we have volatility in our interest rate model. We know that today's interest rates are not tomorrow's and that, as Don said, expectations are not going to line up with where we are today.

So, we create a lattice or a tree, and Don showed you an example of a very simplified tree, where we basically calculate a set of stochastic scenarios. We simulate. We have some kind of a process. And then what we do is we look at each one of the nodes. We look at what the value of the embedded option is at each one of the nodes. And we discount it back using the risk-free rate plus something else. Plus something else is the OAS. And we calibrate that such that it's equal to the starting market value. So, the cost of the option is the difference between the OAS minus the z spread, and the difference between the two is that we've added volatility. Now, note that the volatility assumption is quite critical. If you have a volatility assumption of 10 percent versus 20 percent, the more volatility you have, the greater the value of the option is going to be. So, you want to make sure you understand and have a good volatility assumption in your model to reflect the appropriate value of the option to then obviously reflect the appropriate value of the asset.

Now we have bonds with embedded options, and I did not show a tree/lattice method because Don really covered it. So, he saved me from building a tree. But essentially, what the lattice does is it goes through every point in the path. It looks at the value of the option and discounts it back and comes up with an option value. The market value of a callable bond is the bond price without the option minus that option value. Now, there's a pro and there's a con here. The pro is that you're getting good precision and that there are models today that will actually do this on a path-wise basis. The downside is that you have runtime issues because you're basically building a lattice and running through this lattice for every bond that has an embedded option in it. There are some runtime tricks in some of the models. I'm not familiar with all of them, but the ones that I am familiar with, you can do things like you can give it different time steps. You can either have a lattice that has monthly time steps or annual time steps, and that will definitely help to cut down on some of the runtime.

The other thing that's really critical here, by the way, is that you want to have the best. If you're interested in effective duration and convexity, which is where I'm going to go next in terms of the presentation, you want to make sure that you have good asset market values, because the appropriate way to calculate a market value for a bond with an embedded option is to calculate it at some level point and then shock it up and shock it down. So this would be the appropriate method if you are concerned about effective duration and convexity. And I'm going to come back to that point.

I'm going to go quickly through market value of securitized assets. Here also we're using what I'm calling a nested stochastic simulation approach, which basically means that we have models today that can calculate the market value of the securitized asset at time zero and use that exact same method at future points. If you don't want to do a single path mortgage-backed securities (MBS) calculation anymore, you don't have to. There are basically two different methods, and the tradeoff is runtime versus precision. One is the tree approach where basically you take your starting yield curve and you convert it into your implied forwards. You price along that path, but you can also give it a number of standard deviations, and the model will probability weight these different paths.

For example, you can say that you want five paths, and it'll do two standard deviations above the current yield curve and two below, and it'll price along those paths, probability weight them and come out with a market value. Or you can use a full Monte Carlo method, in which there are tricks to be able to do things like use variance reduction techniques to cut down on the number of runs. I looked in the Fabozzi book over the weekend, and he seems to think that typically you need 1,000 runs through a full Monte Carlo method in order to get a good market value, just to give you some indication. So imagine if you have a portfolio of several hundred of these, and you want to do this every month. You could come back next year to the Val Act, and still be running last quarter's model.

Just to give you an example, we ran a little test in our shop using securitized assets (Chart 7). You see on the left I have Cusip 1 and Cusip 2, and time, where it's a dash, is times zero. I have time 12, 24 and 36. So we're doing this at future points. And we looked at two different things. We looked at number of iterations, and we also looked at altering the random part of the process to see whether we were able to get any convergence. If you look at the hundred iterations, and you compare Random 1 to Random 2, you can see that we converge to a place where we're getting relatively reasonable 1 versus 2 market values. This tells me that we can probably feel comfortable to stop at around 100 scenarios, and this is more of an art than a science at this point.

So we're going to move now, and let's assume that now we have good market values along our paths, but senior management wants to know what the effective duration and convexity of the portfolio are. In order to calculate effective duration we want to calculate stochastically a number of scenarios at some base level, and then we're going to shock up, and we're going to shock down, and just note that the shock, whether you're going 10 basis points up, 10 down, or 50 up, 50 down, matters, usually not a lot, but it matters. So, you want to be able to look at different shocks and understand how they impact your duration because if you have too wide of a shock, you're going to wind up with convexity in your duration number, which you probably don't want to do.

Here's just a little example to show you as the yield curve is flopping around what can possibly happen to your asset duration. So, typically, with a non-callable bond, what happens if my entire yield curve goes up from one quarter to the next? I can expect that my asset duration is going to go down. If my entire yield curve shifts up, and I have a non-callable bond, it's an inverse relationship; the duration of my asset is going to go down. So, in this particular case from 3/31 to 6/30, if you can remember in your minds, the entire yield curve shifts up. So what happens? The duration of my assets should go down, but I look and I see that from 3/31 to 6/30 I have the wrong relationship. I have asset duration of 5.2 at 3/31, and at 6/30 I have 5.62.

So what happened? We broke it out, and we said, "Okay, let's look at bond versus collateralized mortgage obligation (CMO)," and you can see that the right relationship occurred for the bonds because they're predominantly non-callable. We went from 6.93 to 6.84. Yield curve went up. Duration goes down. But CMOs were affected by the prepayments. Rates went up. So what happened? My prepayments slowed down, and it had a significant impact on my CMO duration. You want to look at the relationships from one period to the next, and then break it down and understand and be able to explain what happened.

Finally, I'm going to move to some interesting work that we've done around duration and convexity to dynamically rebalance. The next generation models, the Moses of the world, the Alphas of the world, can handle path-wise calculations now, being able to do market value of assets as well as liabilities. This means that I can calculate the duration of my portfolio of assets and liabilities at future points, not just at time zero. So, we did an exercise where we looked at the effective duration of the assets versus the liabilities and then looked to see whether or not the difference between the assets and liabilities was within a certain predefined tolerance. If it's too short, then I reinvest in a bucket that gets me longer. If it's too long, I reinvest in a bucket to get me shorter.

We start at time zero. We branch out in scenario 1 to node 1, and we take wherever we are in the project at that experience point on scenario 1. So, for example, if we gave a 10-year rate at time zero, that's 5 percent, and in the up-scenario we now have the short rates at 4 percent and the long rates at 6 percent. For example, we will create a set of stochastic risk-neutral scenarios at that point in time and use those to calculate our effective duration of our existing assets and liabilities. We'll then move along to the next node and do the same thing. So wherever we are in the path, we're calculating a stochastic set of risk neutrals.

Chart 8 has an example of what we did. So, here we have calculated the market value of the assets and the market value of the liabilities. We have an asset duration, and we have a liability duration. So if you look at period 3, you can see we' have an asset duration of 3.07 and a liability duration of 2.53. The difference is 0.94. But I want the match to be two. I can afford to go longer because typically when rates are rising and positively sloping, if I go longer, I''m going to earn more yields. So I want to do that. I solve out algebraically, and I say, "Okay, I can afford to invest –in an asset bucket that has a duration as high as 9.14 with ––the cash available before purchases." What you can see in this particular example is 19,828 in period 3.

Well, then I have to look at my predefined asset buckets and figure out whether I have an asset bucket that I can actually invest in with such a long duration. I make the choice to reinvest, and I go to the next period. Now I come to period 6, and I again look at my asset duration minus my liability duration, and now I''m getting closer. I' have a duration mismatch now of 1.48. So, I solve out algebraically again, and I say, "Okay, –I' have an effective duration now that I can afford to invest in of 7.32," and I again find the appropriate bucket that I''ve predefined. I keep going, and you see that eventually I get to period 12, and now I''m a little too long. I' have a mismatch of 2.13, so I switch to a shorter bucket. I' have an effective duration of 3.62, so I go shorter.

This basically allowed us to take effective duration and use it to make decisions about how to reinvest in various different buckets based on looking at future points, seeing where my mismatch is and where I want to be in terms of my mismatch. It also will look at the relationship of the yield curve. So in this particular case it was always positively sloping, but if it was inverted, then I don''t want a mismatch of two. So the mismatch would change relative to the shape of the yield curve and would culminate in also being able to look at alternative GAAP earnings and what the GAAP earnings projections would be based on the reinvestment dynamically.

'**FROM THE FLOOR:** –'My question is on deflators. When we' have multiple indices that we''re modeling, do we run into conflicts trying to back into an appropriate deflator –for a scenario?

MR. WILSON: –I''m not sure whether this is an answer to the question or not, but if it isn''t, you can ask it again. The deflators are independent of the cash flows. Deflators are used to value the cash flows by discounting, and they are independent of the nature of the cash flow. So, –for every simulation coming out of your stochastic generator, like you get short interest rates and long interest rates and equity returns, you get one stream of deflators, and you apply those to any cash flow, and they will give you a value. So, there''s no conflict over what you''re applying them to. '–

–The actuaries first started to look at valuing guarantees they were working in a real-world sort of simulations and were tending to underestimate the cost of the guarantees. It was only when the risk-neutral type approach was used–. We brought in the market–consistent approach, and by market–consistent I mean what the Wall Street people were using it to value assets. But we got to a methodology that put a more appropriate value on the guarantees. '–'–'–The reason why 'a valid –price on the guarantee is that you could consider perfectly hedging a guarantee against an option bought in the market. How much would that option cost you assuming you could buy it? It would –be valued by the Wall Street types, using their methodology. So, –if we apply the same methodology on the liabilities, then we get at least to the price of a perfect hedge, if that''s available. Of course it isn''t generally. VA guarantees tend to be 20 to 30 years out. You can''t buy hedging assets beyond, say, 10 years. –By using the same methodology we can be sure we have' a reasonable price.

MS. COOPER: I agree that there are a lot of questions. –There are a lot of questions about how you determine risk-neutral versus real world. If I look at my cash flows under real world or if I look at the value of my guarantee in the real world versus risk-neutral, how do I sort of equate the two? I think the bottom line is that the price is the price, and it''s the way that these are priced in Optionland and in Wall Streetland, and so to be able to do derivatives-based hedging this is how we price in a risk-neutral context to equate them to the appropriate derivatives in the asset world. I''ll just add that for long-tail liabilities you actually can have a hedge that works and balances. It''s just that it has to be a dynamic hedge. It can''t be a static portfolio.

FROM THE FLOOR: I have a question for Don. You mentioned that if you' have 20 or 30 funds in a product, that you don''t want to model each of them with its own assumption, but you want to map it to maybe three to five indices or funds. What advice do you have? How should one go about figuring out how to map that?

MS. COOPER: I''ll take that. T–his is one of the things I do every day. There are a number of ways you can do it. You can do it the quick and dirty way. You can take your funds and look them up in Morningstar and see what the categories are and do a quick mapping that way, or you can do something much more rigorous. Depending on the use of what you''re doing, I''d recommend different things for different purposes. –You can create multiple regressions, historical regressions, of the fund against the various different indices, and again if you''re going to hedge. For example, you want to look at the multiple regressions against those indices that you would potentially be hedging against. –So, S&P 500, NASDAQ, Russell, and basically you would have a calculation, a formula that would be alpha plus beta 1 of S&P plus beta 2 of Russell plus some epsilon, and that''s how you would do it. You would need to look at those and periodically update them, quarterly or at least annually, because the behavior of the funds will change as the equity markets change. So you don''t want to just leave those, and it is really important.

FROM THE FLOOR: Just a quick follow-up. I would imagine that if that fund had been around 20 years versus one, you''d have a greater amount of comfort in the historical mapping. At what point would this occur? What''s the rule-of-thumb?—

MS. COOPER: I would say it depends because sometimes what happens with the historical mapping is it moves over time. So, a fund that might be a growth fund might –behave more like a small mid-cap fund. In a certain period of time it may behave more like an S&P 500 over another point in time. So you just need to understand the trend of the fund versus what''s happening in the real world.

Chart 1

Economically coherent

- Efficient markets
- What happens next?

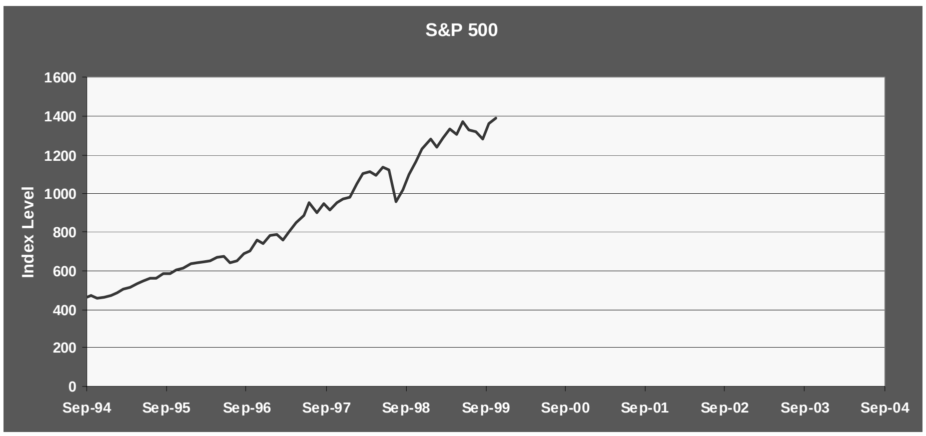

6

Chart 2

Economically coherent

- Full term structure
- Prices now for a cash flow in the future
- Required for correct pricing of future liabilities

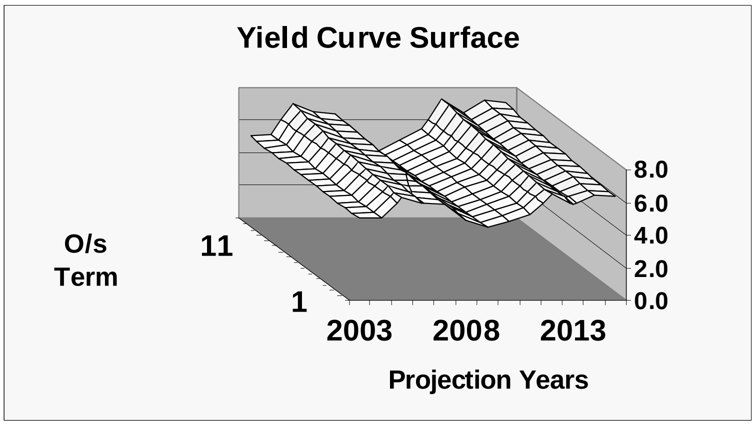

11

Chart 3

Market Conditions

Term Structure of U.S. Yield Curve

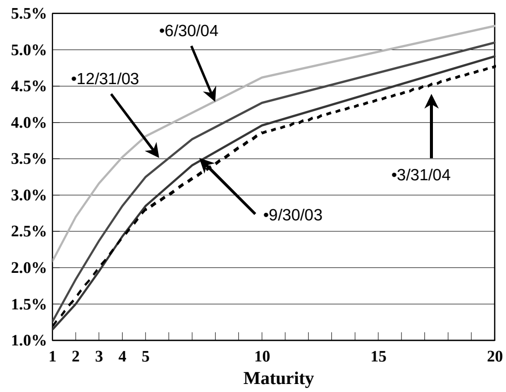

	9/30/03	12/31/03	3/31/04	6/30/04
1 Year	1.15%	1.26%	1.20%	2.09%
2 Year	1.50%	1.84%	1.60%	2.70%
5 Year	2.85%	3.25%	2.80%	3.81%
10 Year	3.96%	4.27%	3.86%	4.62%
20 Year	4.91%	5.10%	4.77%	5.33%

Chart 4

Market Conditions

Credit Spreads of U.S. Corporate Bonds

Average Credit Spreads

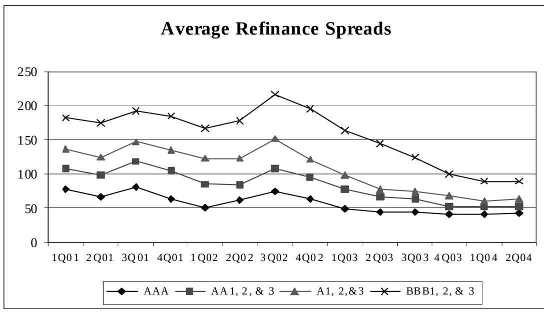

Grade	3Q03	4Q03	1Q04	2Q04
AAA	44	41	41	42
AA1, 2, & 3	63	52	51	52
A1, 2,&3	74	68	61	64
BBB1, 2, & 3	125	100	89	89

Credit spreads continued to slightly decline and have converged to a narrow range.

Chart 5

Current Credit Environment – Falling Recoveries

Avg. Annual Speculative-Grade (Original Issue) Defaulted Bond Recovery Rates, 1982-2002
Source: Moody's annual default and recovery study (2/2003)

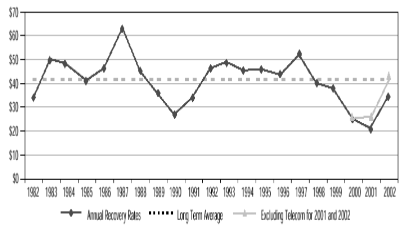

Recovery Rates and Macroeconomic Conditions
Source: Altman & Brady (2001)

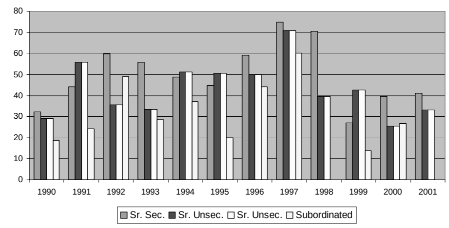

- Record defaults have contributed to declining recoveries (supply and demand)
- Seniority continues to be the strongest indicator of recovery potential; subordinated debt holders have suffered the most through this credit cycle
- Combination of record high defaults and weak recoveries have resulted in large credit losses for many investors

Chart 6

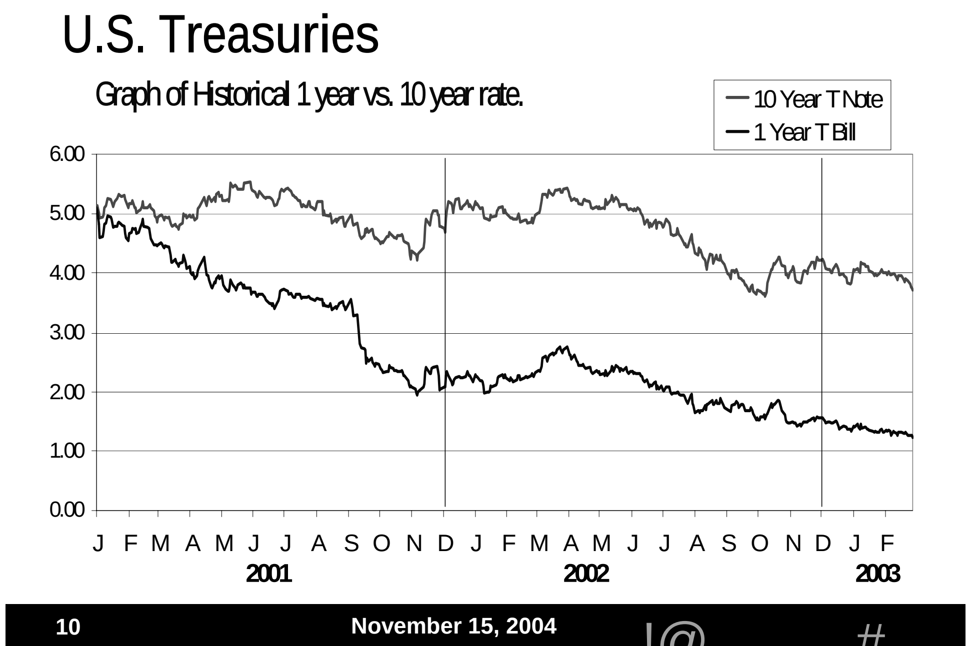
U.S. Treasuries
Graph of Historical 1 year vs. 10 year rate.
10 Year T Note
1 Year T Bill
6.00
5.00
4.00
3.00
2.00
1.00
0.00
J F M A M J J A S O N D J F M A M J J A S O N D J F
2001
2002
2003
10
November 15, 2004

Chart 7

Market Value of Securitized Assets – How many Scenarios are enough?

Number of Monte Carlo Iterations

		Random Seed 1				Random Seed 2			
Cusip	Time	1	10	50	100	1	10	50	100
1	-	5,933,117	7,127,458	7,147,145	7,000,619	6,171,104	6,950,740	6,931,834	6,948,469
1	12	5,318,772	6,125,019	5,725,299	5,792,399	6,337,256	5,756,095	5,822,755	5,850,955
1	24	4,886,750	5,262,374	5,183,868	5,082,375	4,629,070	5,084,450	5,145,372	5,121,315
1	36	4,546,624	4,523,881	4,439,146	4,511,356	3,966,068	4,600,717	4,557,651	4,546,852

Cusip	Time	1	10	50	100	1	10	50	100
2	-	27,369,283	32,122,142	32,750,082	32,181,847	28,404,583	32,137,585	32,142,234	32,102,684
2	12	31,184,051	33,952,955	32,348,788	32,510,908	34,531,282	32,791,590	32,673,809	32,650,089
2	24	32,859,168	33,123,940	32,962,499	32,862,102	30,811,663	33,026,421	32,921,754	32,959,415
2	36	30,444,355	30,649,155	30,246,076	30,471,952	29,263,658	30,593,718	30,418,251	30,435,805

Chart 8

Dynamic Rebalancing to A Target Duration

Dynamic Reinvestment Strategy	0	3	6	9	12	18
Total Market Value of Asset (incl. Cash)	91,482	105,008	122,016	139,768	158,040	194,748
Total Market Value of Liabilities	87,353	103,092	119,071	135,148	151,347	183,387
Market Surplus	4,129	1,915	2,945	4,620	6,694	11,361
Cash Available before Purchases	-	19,828	19,177	20,231	20,052	39,450
Market Value before Reinvestment	91,482	86,006	102,947	119,993	138,310	155,778
Asset Duration	2.91	3.07	4.49	5.01	5.29	4.82
Liability Market Value - Base	87,353	103,092	119,071	135,148	151,347	183,387
Liability Duration	2.50	2.53	2.54	2.53	2.49	2.43
Asset/Liability Duration Mismatch	0.94	0.94	1.48	1.84	2.13	1.90
Acceptable Duration Mismatch	2.00	2.00	2.00	2.00	2.00	2.00
Target Effective Duration for Reinvestment	-	9.14	7.32	5.47	3.62	4.81