RECORD OF SOCIETY OF ACTUARIES
1995 VOL. 21 NO. 1

ADVANCED CASH-FLOW TECHNIQUES IN PRODUCT PRICING

Moderator: CHERYL A. KRUEGER
Panelists: DOUGLAS A. GEORGE
CHERYL A. KRUEGER
Recorder: CHERYL A. KRUEGER

The session will focus on the issues of dealing with management and applying cash-flow techniques to the process of bringing profitable products to the market.

MS. CHERYL A. KRUEGER: Doug George and I will talk about advanced cash-flow techniques and product pricing. I've been with Tillinghast in Chicago for almost five years, where I have been involved with cash-flow testing, embedded-value calculations, asset/liability management, and other life and annuity projects. Before that I did pricing at a couple of companies—one was mainly an annuity company, and the other had a diverse portfolio of individual life insurance and annuity products. Doug is with Avon Consulting Group. He has experience in product pricing, cash-flow testing, asset/liability management, mergers and acquisitions, and other strategic financial analysis. In terms of asset/liability management work, he has experience in scenario-based and option-adjusted approaches to risk management.

I will discuss why cash-flow methods can and should be used in product pricing. I'll present a couple of techniques you can use including just a simple spread adjustment to your interest-sensitive product pricing using scenario-based analysis in product pricing. Doug will talk about system implications and option pricing.

I'll start off with a simple summary of what cash-flow testing is. Basically, it's the process of determining whether the necessary cash will be available to pay out benefits in the future. It involves projecting and analyzing the timing and amount of future asset and policy-related cash flows, so it's not just a liability-based analysis. Of course, there are many different types of cash flows; some of those cash flows will be interest sensitive, including surrenders and the sale and early maturity of assets. In terms of pricing, interest-sensitive characteristics of products were traditionally ignored until a few years ago. Now I think that more and more companies are taking them into consideration.

Cash-flow testing used on in-force blocks has become common because of legislation and because rating agencies have become more knowledgeable about cash-flow testing. Also, what we've learned from cash-flow testing can be applied to management to make it more than just a regulatory exercise. What have we learned?

First, the level interest assumption is a dangerous thing. Also, a counterintuitive result may not be. In other words, you can't just pull the results out of the computer and assume that they're right. In addition, the process itself is beneficial to management. We'll get into that a little bit later on when we talk about how the changes to the views of pricing can be communicated to management.

One thing that is very important in terms of pricing is how we value the policyholder options that are in the product. We're specifically talking about interest-sensitive products and the options therein. There is an opportunity for policyholder antiselection in most of our contracts; most interest-sensitive contracts allow surrender at book value, and the assets

that need to be sold to cover cash flow (if they do need to be sold) are sold at market. I think now more and more insurers are charging for this right; determining how much to charge is a pricing problem.

The duration of the liabilities can't be matched with the duration of assets where there are book value guarantees, and that's a valuation actuary problem and a pricing problem. And the strategic risk/reward trade-offs can best be made after risk/reward management or after the measurement of risks and rewards in the products.

Obviously, part of the problem is that it's difficult to match the durations of the assets and the liabilities. When interest rates are low, the liabilities have a long duration and assets have a short duration; as interest rates get higher, that relationship switches around. Thus there is a difficult management problem on an in-force block and in pricing.

Why should we use cash-flow testing in pricing? First, we can measure the true cost of the policyholder options rather than just using a level interest environment and assuming a set surrender rate. We can evaluate alternative investment and interest crediting strategies. Maybe we're using a portfolio less a spread for a crediting strategy and there may be some situations where we might want to take other things into consideration—for example, competitor rates—to set credited rates.

Also, we can explore management's appetite for risk. One of the benefits of this analysis is the communication that occurs as a part of the process. There are issues other than just the actuarial techniques that need to be addressed; so certainly management needs to look at how the analysis is done, what the parameters are, and how the company will measure risk and return.

It reduces the likelihood of surprises. It also helps to evaluate the surplus impact of a product and its associated risk. All these things together help cash-flow testing methods in pricing tell a better story.

There are some limitations to cash-flow testing. One of them is the assumptions that go into cash-flow testing. There are quite a few assumptions in terms of reinvestment strategies, policyholder behavior and others. It's also difficult to capture the value that's added by investment expertise—some companies feel they get much of the benefit out of their investment departments. That is difficult to capture in this type of analysis.

It's also difficult to project active asset management strategies. As I mentioned before, the assumptions in terms of policyholder behavior are difficult to set. We have recently been in a period of some rising interest rates, but policyholder behavior hasn't always been what we projected that it would be in our models.

Tillinghast did a survey a couple of years ago to find out how many companies are actually using cash-flow testing methods in product pricing. We haven't updated the survey, so I thought I would take a quick informal survey here. First, how many pricing actuaries do we have in the audience? There's a fairly good number—almost everybody responded. How many of you were using cash-flow testing techniques two years ago in your pricing? It looks like four or five. How many are using it now? There is quite a significant number (maybe 20 or 25) that are using it now. Not surprisingly, the survey

results are out of date. There have been many changes since this survey was taken, but I will go ahead and review what people were doing several years ago.

We received responses from 100 companies. Results are shown in Charts 1–3. Sixty percent of the companies weren't using cash-flow testing methods in pricing. Of the remainder, a fairly large number were doing some stochastic analysis; others were doing some option-pricing analysis. Many people were using cash-flow testing methods to test the investment strategies. Many were setting the interest crediting strategy. Some are setting the asset duration target for communicating the investment requirements for new products to the investment department. A few companies were using cash-flow testing techniques to value the options in a product, such as bail-outs or market-value adjustments.

CHART 1
USES OF CASH-FLOW TESTING IN PRODUCT PRICING

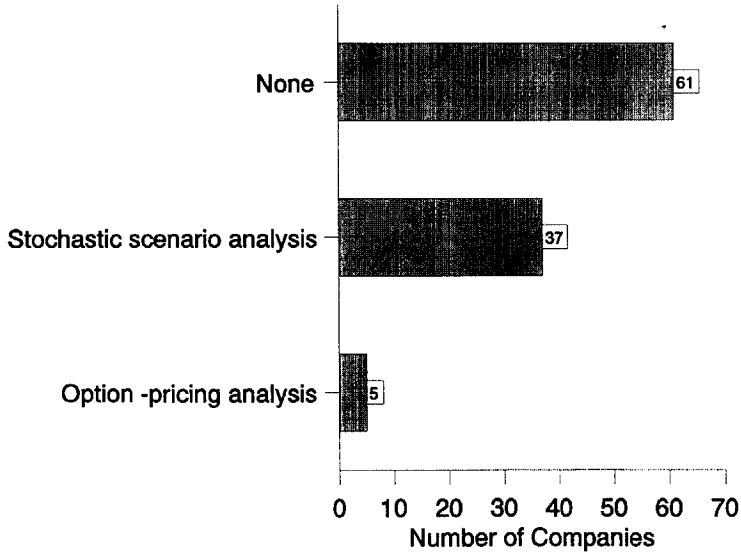

There were only a few companies at that time using option-pricing analysis. I would guess that there are quite a few more now. Again, testing the investment strategies was the most prevalent use of that technique.

I will talk a little more about multiscenario cash-flow testing (stochastic scenario analysis). Doug and I will talk about three different approaches to incorporate interest sensitivity into your pricing. I will talk about a simple method to use if you're not doing anything else—just making an adjustment for interest rate risk in terms of a spread adjustment to your pricing. Second, I'll talk about asset/liability efficient frontier analysis and how it can be used in pricing. Finally, Doug will talk about option pricing.

CHART 2
STOCHASTIC SCENARIO ANALYSIS

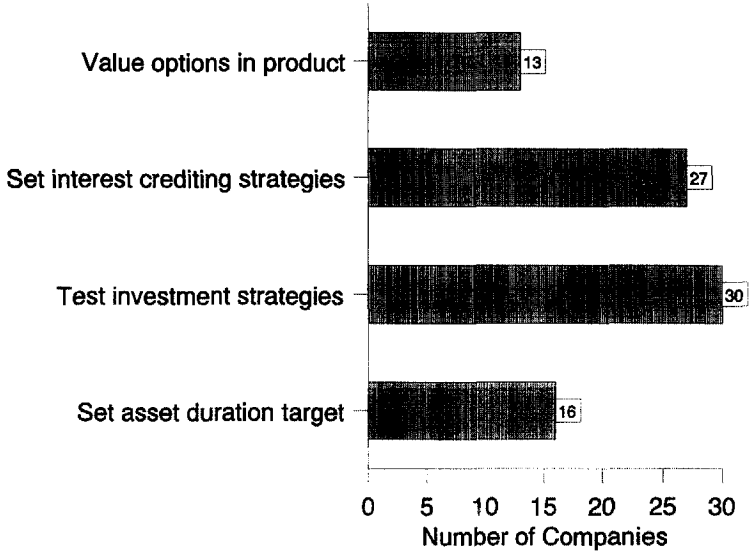

CHART 3
OPTION-PRICING ANALYSIS

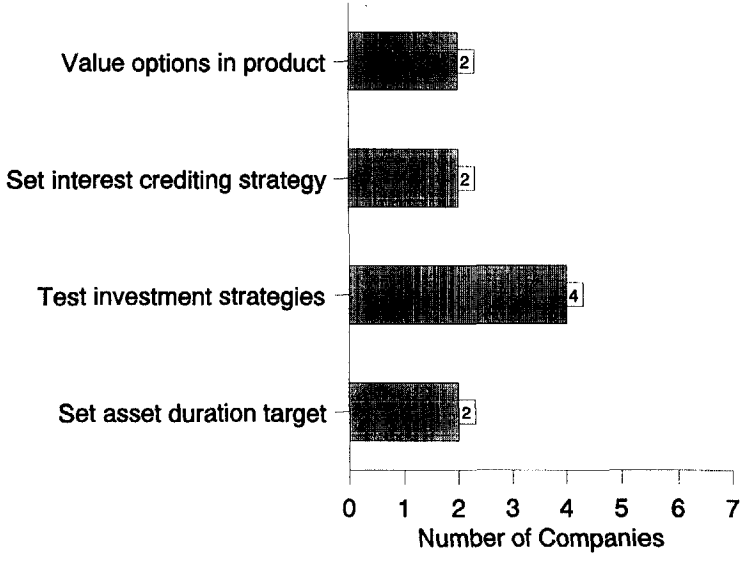

In multiscenario cash-flow testing, the result that we get is presented with respect to company profitability or product profitability. We define the cost of the options in terms of company profitability. The range of results allows for some understanding of potential downside risk. In other words, things that we talked about before—what happens to policyholder options when interest rates change—are actually measured here, rather than just saying "We know they're out there."

When you're pricing, it is important to determine which method you will use. One important consideration is, what question are you trying to answer? Expected profitability of a product, downside risk, and desirable asset portfolio attributes would be some factors to consider when you're determining which asset/liability management technique you will use in pricing.

Are answers in a form consistent with corporate goals, and are they understood by senior management? I think one thing that will help determine whether your approach will be understood by senior management is what is being used on the in-force block. What are they familiar with, what are they comfortable with, and what types of analyses are being done that you can then use in pricing? Many of the decisions that need to be made in terms of this analysis have probably already been made and should be reviewed in deciding how pricing will be done.

Multiscenario cash-flow testing will result in a range of profitability measures. Typically, the average results over a set of scenarios (such as the average present value of profits at your hurdle rate) are less than the results for a level interest scenario. That is almost always true because the policyholder options embedded in the product that are now being valued weren't being valued in the level interest environment.

One technique that can be used is fairly simple and straightforward. If you're not doing anything else in terms of valuing options, this might be used to get a feel for how much cost there is to the options that you offer to your policyholder.

Doug and I have worked out an example of each of the techniques that we will talk about. We used a basic single-premium deferred annuity (SPDA) product with a typical surrender charge that grades from 7% of fund value to zero. That's the basis for all the numbers that I'm using. I haven't presented any details here about the products because we want to try to focus on the results and how they can be used, rather than the specific technical details of the product.

These are the steps that you would take for a spread adjustment. First, you would evaluate your product under a level interest scenario which you typically would do in pricing. Then you would subject the product to some stochastic interest rates and include some assumptions in terms of reinvestment strategies and policyholder behavior in different interest rate environments. After the second step, you will end up with a profit level that's lower than you ended up with in the first step. To get it back to the original profit level, you increase the spread that you're using to calculate your credited rate. Then you go back to your level interest rate scenario model and you include that spread as another cost to the company. Essentially, you are calculating an expense or an additional spread that you will need to adjust for the risks of policyholder behavior and interest rate sensitivity.

For example, start out by profit testing the SPDA product in a level interest rate scenario. We started with a spread of 145 basis points, which gave a 14.2% return on investment (ROI). We then took the same product and ran 100 interest rate scenarios in an asset/liability model. We came up with an average rate of return of 12.9% (down from 14.2%). I'll qualify the actual number that we're coming up with in terms of spread and tell you that we had a fairly short investment strategy here. We didn't have any bonds longer than ten years, so the result of five basis points isn't a rule of thumb. It's certainly something that would vary depending on the investment strategy and the product as well.

Next, we put in a new spread of 150 basis points in the asset/liability stocahstic model, and the average return is up again to 14.4%. So for this product (investment strategy) and set of assumptions, our C-3 interest rate risk cost is five basis points.

Now, as long as we don't make too many changes, we can go back to our level interest rate model and all we need to do is include the additional five basis points as a cost. We were working with a spread of 145 basis points and now we'll go down to 140 basis points. Our return will go down. We now adjust our product so that we get back up to the previous ROI level. This technique is something very simple, but it gives you a feel for what the cost of those options are.

A more elaborate and sophisticated method to use for multiscenario analysis is the asset/liability efficient frontier analysis. It's a means of evaluating various strategies, in terms of reward and risk. The definitions of risk and reward reflect management philosophy. How these things will be communicated with management might depend on what you're using on your in-force block, because one of the big decisions to be made in terms of doing efficient frontier analysis is, how do we measure risk in the company, and how do we measure returns of the company? What's important to us and how will we measure it?

One thing about the asset/liability efficient frontier technique is that it's a fairly straightforward communications tool. Once you have the risk and returns defined, you're looking at different options, different strategies, or different product features. You price each of those strategies under multiple stochastic scenarios and you take the average results measured in terms of your definitions of risk and reward and plot results on a graph. At that point you can say, for example, here's a product with one crediting strategy and here's another crediting strategy, and which one is better? Obviously, this one will give us a better return with less risk than this one.

Once you get used to this approach you can look at a graph like Chart 4 and assume that anything that's on this efficient frontier line will be where the choices are. For example, we ask, "which is better B or C?" It depends on the company's philosophy in terms of risk. It may prefer to have less risk and less return or it may prefer to have more risk and more return. That's where the decisions need to be made because you can eliminate all these other strategies that are off the line. For the points along this line, you need to decide how much risk the company wants to take on. Again, if your company is using this type of approach on its in-force block to make decisions, much of the work in terms of defining risk and return is already done; much of the education process on how we look at risks and returns has been done. Then the questions for pricing are, where do we make our choices, and what options do we look at?

CHART 4
STRATEGIES ON THE "EFFICIENT FRONTIER" ARE SUPERIOR

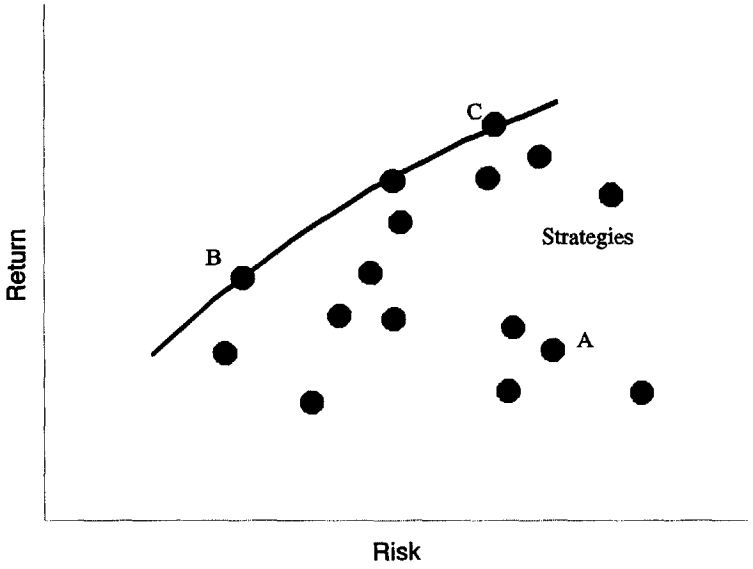

To perform the efficient frontier analysis, first, choose the risk and return criteria. These should be company-wide objectives, not department objectives. That is real important in terms of relating the results to senior management. If you're using the same criteria on the in-force block, it is easier to get some buy in on the results. The setting of risk and return criteria may be the most difficult part of the process. I've been involved in efficient frontier analysis with a couple of companies. Some people are concerned about GAAP profits; some people are concerned about the level of surplus; mutual companies may be concerned about accumulated surplus and/or dividends. Determining what's important to management is a difficult but very important part of the process. If the criteria in the analysis aren't what the company is actually looking for, then the analysis isn't going to be successful.

The next thing you do is choose the product features to be compared. For example, if you're looking at crediting strategies, chose a few to be tested. Run multiple scenario models using each of the options that you want to look at. Compare the results by putting them into a graph; then you can communicate them.

Again, much of the information you use may already be put together because of the cash-flow testing exercise, and it mainly involves coming up with a new business model.

In an example of using the efficient frontier analysis to compare some crediting strategies, I ran results using five different types of crediting strategies. I called the first one "unconstrained portfolio rate less spread." No matter what assets are earning and no matter what

the competition looks like, the credited rate is simply the asset earnings rate less the spread.

The second one is the portfolio rate less a spread, but with competitive constraints. For example, we may assume we will not be more than 50 or 75 basis points out of the market.

The third one is the same as the second one except we will never be so far out of the market, and we will never earn less than 75 basis points, either. So we've got a little more constraint in our crediting.

The next two strategies are based on competitor constraints only. We ignore the asset earnings rates altogether and follow the competitor in strategy four. In strategy five, we follow the competitor halfway up or down, so we start at the competitor rate and then follow the competitor rate halfway whenever the rates have moved.

I used a couple of different return criteria. In the first set of results, shown in Chart 5, the risk was measured as the probability that our internal rate of return falls below 11%. We measure return as the present value of distributable earnings at a 15% hurdle rate (the median of the 100 scenarios).

CHART 5
EFFICIENT FRONTIER

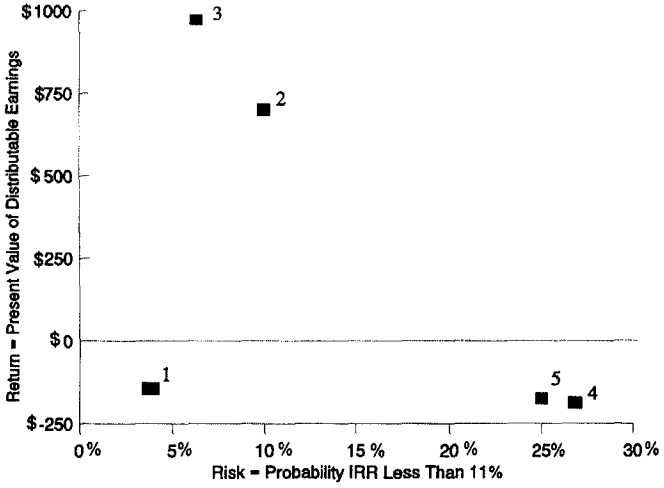

In terms of these crediting strategies and these risk and return definitions, we can see that the management choices here are basically clear, that the decision would be to go with a portfolio rate, less a spread, but have some competitor constraints and also make sure that we are getting at least a 75-basis-point-spread. Understand that point 3 has the highest return at the lowest risk on this graph. Another possibility is to go with the less risky option of point 1, which doesn't meet the profit criteria.

FROM THE FLOOR: And these all have the same assets?

MS. KRUEGER: Yes, everything is the same except for the crediting strategy.

Now I change the risk definition to be the standard deviation of the results at a 15% hurdle rate. My risk now is measured in terms of variability of results rather than level of results. Again, I use the same return definition as in the first set of projections.

Now I have a slightly different graph, as shown in Chart 6. I think that the only place where I would make a different decision would be if there were certain strategies that I had eliminated. The position of the options in relationship to each other have changed even though the decisions haven't changed that much. However, assume that you were choosing between strategies 2 and 5, for example. Using the first risk criteria, we clearly would pick 2. Under the second risk criteria, we might consider the less risky option of 5, where under the first risk definition, number 5 is clearly eliminated from consideration. So you can see that when we change the definition of risk, it changes the results. In this case it's not a severe change, but the point is, there can be a change depending on what the risk and return definitions are.

CHART 6
EFFICIENT FRONTIER

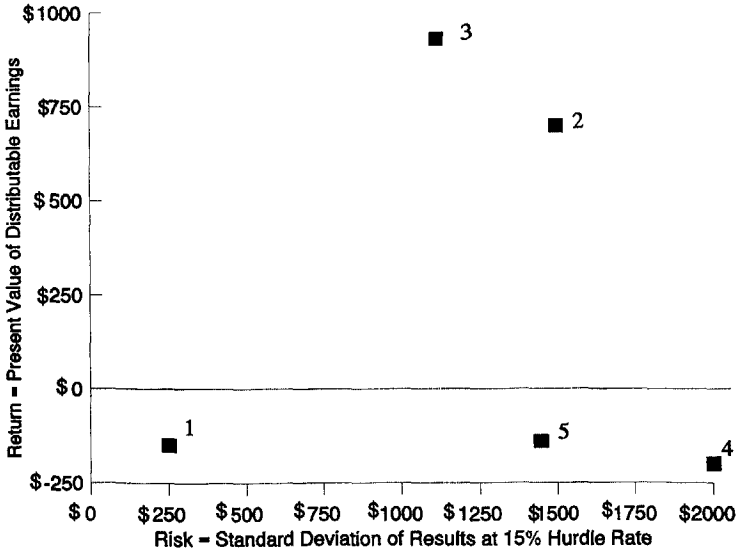

That covers the two methods that I was going to talk about. I also want to talk a little bit about the advantages and the disadvantages of the multiscenario approach. The options are valued in terms of the impact on company profitability.

We look at a range of results, not just one result. So we can look at downside risk analysis. The efficient frontier customizes the risk and reward definition. The company gets to choose how it will measure the results. It also reflects the cost of capital because we can include target surplus in the analysis and we can use company specific discount rates when we're making the analysis.

The disadvantage is that in order to get some meaningful distribution tales on the down-side risk, you do have to run many scenarios. The output from the projections can be overwhelming unless they're summarized, possibly by eliminating obvious nonchoices, or looking at something simple like the efficient frontier graph rather than looking at a set of 100 outputs from each of the scenarios.

That's all I have on the multiscenario approach. Now I'll turn it over to Doug. When he is finished, we will have time for questions.

MR. DOUGLAS A. GEORGE: I will talk to you about option-pricing approaches to asset/liability risk pricing which is in contrast to the scenario analysis that Cheryl showed. Scenario analysis can show you the effects of the options in your assets and in your liabilities. Option pricing is a technique where you explicitly value those options. It's done usually by projecting a number of interest rate paths over a future interest rate space. You project the cash flows of an asset or a liability along each of those paths, take a present value along each path and then probability weight the present values across all the paths to determine a market value. On assets, it's a real market value. For liabilities, it's a theoretical market value. I believe you should do this kind of analysis in conjunction with the scenario analysis that Cheryl was talking about. There are pluses and minuses to each type. Option pricing is a good way to show you the approaches to change your asset/liability risk, to fix it, if you will, or change the risk profile, whereas scenario analysis gives you the effects of those changes on your earnings.

I want to expand my talk a little bit to include in-force pricing as well as new business pricing. Before I go into the techniques used to perform option pricing, I want to tell you a little bit about the systems and processes that I use to get this accomplished.

I have a few concerns about using my actuarial software to perform asset/liability risk analysis. First is the interest rate process. In most software packages, the option-pricing process is rather simple. I'm not sure if it really captures the nuances of the options in my assets and my liabilities. Second, asset modeling is often limited. Collateralized mortgage obligations (CMOs) can be a problem. My actuarial software doesn't give me the deal structures that I need to model different types of CMO tranches. Prepayment models are also usually fairly basic. There's often no database to provide updated collateral and prepayment speed information.

Finally, asset model creation can be difficult because the packages aren't often designed to accept the information well from the asset administrative systems. Computer run time can be an issue especially if my software is built in APL, which is rather slow. Finally, there's a consistency concern. If I build my assets and my actuarial software, the analysis is not directly comparable to the analysis that my investment people are doing when they're trading the assets, so if I get a duration of four and my asset people tell me the duration is 3.5, what do I do? If my actuarial software is not consistent with what my investment people tell me, my results will not be as credible.

The method that I use to get by these concerns is systems integration, as illustrated in Chart 7. I take my actuarial software and link it up with the investment software that my investment people are using to trade their assets. This can be a homegrown system developed internally, or it can be one of the commercial packages that's available. You can do this, for example, with software from global advance technology (GAT), capital management sciences (CMS), or other types of commercial or homegrown software.

CHART 7
SYSTEMS INTEGRATION

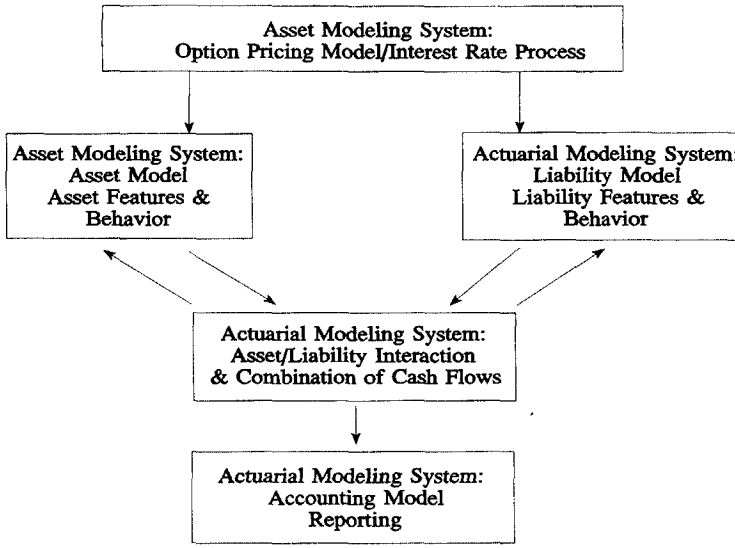

I usually start with my option-pricing model and I usually take that from my asset software. I usually use one of the commercial packages because they have been thoroughly tested and have very robust option-pricing models. Their option-pricing models have been proven to price both short and long options rather well.

I connect that to my asset modeling software. This can be the same system or it might be a different system. It might be your own homegrown system that your investment people have developed to project cash flows for your assets. I also connect it to my actuarial modeling software so I can price my liabilities using the same option-pricing model that I've used for my assets.

Next I need to have a method of interacting between my asset and liability models. I usually use my actuarial software to perform the asset/liability interaction. I do this because the actuarial software has capabilities for projecting future cash flows—determining future cash flow needs at a point in time along projections or along interest rate paths or scenarios. They allow me the ability to buy and sell securities in the future and to perform cash-flow testing analysis, whereas the investment software packages don't typically offer this. Although I have used the investment modeling software for the

interaction piece, usually I'm limiting myself to only option-pricing analysis, rather than scenario-based analysis.

Finally, I need a method for accounting and for reporting the results. Once again, I go to my actuarial software because it has the statutory and the GAAP accounting features in place. I've got the ability to measure required surplus or benchmark surplus or risk-based capital needs. Finally, for the reporting you can really use either system, but it's natural to use the same system that you finished with coming from the accounting model, so I usually use the actuarial system there.

There are some concerns with a system such as this. There's the timing of the cash flows. Sometimes my asset software might have monthly cash flows, my actuarial software might have quarterly cash flows. There's the interest rate definition. My asset software might have a bond equivalent yield curve. My actuarial software might have a spot curve definition. And finally, there are default modeling needs. Asset modeling software doesn't typically allow you to explicitly price defaults or C-1 risk, so we need to incorporate those. In practice I've found all these concerns can be resolved. Timing of cash flows can be brought up so that they are consistent. The interest rate definitions can be converted so that they are consistent and the default models can be added to the asset modeling software to make this system whole.

I see a number of advantages to my integrated system. First, I'm getting the best analytics for both sides of the balance sheet. My assets are modeled in the software that best models them, as are my liabilities. The interest rate process is also the best one available and it's consistent. If I'm going to calculate a duration of an asset or a liability and compare the two, I'm using the same process or the same model to calculate that duration so that I can be sure I have an apples-to-apples comparison.

Model creation and maintenance is minimal. Although still there's a lot of work involved, it's less than is typically needed to recreate asset models in actuarial software and vice versa. The system is flexible and in a minute I'll discuss the different uses for a system like this.

It's also consistent with the independent analysis that my investment people are doing so that they will "buy into" the results that are produced. I will not have to fight them politically to get them to buy into my asset/liability results which might happen if my model were different than theirs.

Finally, computer runtime can be saved. Using two systems, the assets and liabilities can run simultaneously as opposed to back to back. Also, asset modeling software is usually written in compiled code so that you can get much better run time than in APL.

There are a number of analyses you can do with the integrated system. You can do financial forecasting, capital planning, and of course, cash-flow testing. You can tie your product pricing to your corporate modeling so that you can develop a feedback loop to give you better pricing assumptions based on your in-force experience. Of course, you can do your product pricing including your asset/liability risk analysis, and that's what we will talk about.

I'm will demonstrate an SPDA similar to the one Cheryl used for assumptions that are similar, but that may not be exactly the same. My purpose here is to demonstrate the different techniques of option pricing and the pluses and minuses of each; I will not provide actual answers.

First, there's option-adjusted spread (OAS) analysis as illustrated in Table 1. This is a technique where we take the current yield curve and perform asset/liability option pricing. We use the all-in-cost cash flows, including benefits, expenses, and commissions.

TABLE 1
OAS ANALYSIS—SPDA

Run	OAS	Static Spread	Option Cost
Base	68 basis points	48 basis points	20 basis points
Increases Initial Rate 50 basis points	88	66	22
—Difference	20	18	2
Reduce SC 1% per year	81	56	25
—Difference	13	8	5

The first thing we do is project the static spread and that's in our middle column. For our base SPDA, the static spread is 48 basis points. This is determined by projecting the cash flows along the forward rates implied by the current yield curve and discounting back at a spread above treasuries, solving for the spread that equates the present value of the cash flows with the sales price of the annuity. When we project the cash flows and discount back to the market value sales price, we find that we need 48 basis points above the treasury curve to equate the two values.

Next I extend this to do option-adjusted spread analysis and that's illustrated in the left column. Here I encompass a full range of future interest rates. Rather than just projecting along the current yield curve, we're also going to encompass up and down paths going forward. When I include my up and down paths, I project my cash flows along all the paths, discount backward, and I find that it takes 68 basis points of spread above treasuries to equate the present value of those cash flows to my initial premium. This is due to interest-sensitive behavior. On some of the paths, the cash flows will occur sooner than under the static path so that when I weigh all the paths together, the value of the cash flows increases and therefore the OAS goes up as well.

What does this all tell me? These numbers are nice, but what do you do with them? I know that in order to break even on my assets, I need to earn an OAS of 68 basis points averaged across all interest rate scenarios. Second, I can calculate my option cost, which is the difference of the OAS spread versus static spread and I find that to be 20 basis points.

This is the cost of the options that are in my SPDA. If interest rates weren't going to go up or down, and if we could continue to just project that interest rates did follow the static curve, and we were sure of it, we would save 20 basis points of cost.

I can also test asset/liability strategies. If the OAS on my liabilities is 68 basis points, this is what I'm paying my policyholder to get that money. I need to buy assets that are earning higher than a 68-basis-point OAS in order to make a profit on an economic basis. I can test asset/lability strategies that maximize the difference between the OAS that I'm earning on my assets and the OAS that I'm paying on my liabilities.

Finally, I can explicitly price the options in my liabilities. Table 1 shows that if we increase the initial credited rate by 50 basis points, that costs me 20 basis points of OAS. If we reduce the surrender charge 1% per year, that costs me an extra 13 basis points of OAS.

If my marketing people come at me and tell me that my product is doing poorly and we need to do something to beef it up, I know we have two choices. We think they're equally valuable from a marketing standpoint. One will increase the initial creditor rate, and one will reduce the surrender charge 1%. Which one do you think we should do? This technique tells you. It tells you that reducing the surrender charge is less costly averaged across all interest rate scenarios than increasing that initial credited rate, so I will choose reducing the surrender charge.

In a similar manner, if I have two products in my portfolio, and one of them has a reduced surrender charge (1% above another one), then I should probably offer roughly 30 basis points of initial rate higher on the one with the higher surrender charge in order to have equivalent profitability averaged across interest rate scenarios. Since 30 basis points of initial rate is worth about 13 basis points of OAS, the 30 basis point initial rate difference makes me economically indifferent between selling one product versus selling the other. I don't care which one they buy, because the fair value to me is the same.

This is option pricing in its purest form. The real purpose in the design of option pricing is to price assets and liabilities going forward so that we're explicitly pricing the option. We must find out what it's worth.

Table 2 shows a couple more examples for my SPDA. My SPDA has a bailout in it and the bailout goes for five years. It's at say one-half of 1% below the initial credited rate so that the policyholder can withdraw without surrender charges if the credited rate drops below this level during the first five years.

TABLE 2
OAS ANALYSIS—SPDA

Run	OAS	Static Spread	Option Cost
Base	68 basis points	48 basis points	20 basis points
Remove Bailout	58	48	10
—Difference	−10	0	−10
Remove Bailout (in-force)	64	48	16
—Difference	−4	0	−4

Here I test the effect of removing that bailout. When I calculate my static spread, the base run versus removing my bailout, I see that it doesn't change. I've got 48 basis points for both.

This shows you the fallacy, if you will, of static pricing. If you were to just try to price that bailout under the current yield curve going forward, the bailout is never triggered, so the bailout doesn't cost me anything. Obviously, we know better than that. We know that bailouts can be costly even if we're only doing static pricing, but the question is how do you value that bailout, how do you know what it's worth? Option pricing can tell you. When I remove my bailout I see it's worth ten basis points of OAS. Because the static spread has not changed, the option cost is also ten basis points and that's a natural relationship because the bailout is a pure option and has no static cost to it. It is simply a pure option that only costs when interest rate movement occurs.

Finally, I show an in-force product and this one has been in force for three years. You can see that the OAS for the bailout has gone down quite a bit. As we get towards the end of the bailout period, it has much less value, and of course, once we get beyond the bailout period, the SPDA will cost the same as if it did not originally have the bailout.

The OAS analysis is a starting point, but it's limited in that it doesn't give you the specific change in values as interest rates go up or down. We take an average cost across all interest rate scenarios, but what happens if interest rates move up or move down or zigzag? I'm not sure how the values of my options change. For that we can use a technique called price behavior curves.

We can do this because, once again, we have the integrated system. We have an option-pricing model that's connected to my assets and my liabilities to give me option prices for both in a consistent manner. Price behavior curves are a good way to picture the asset/liability risk in your products. For duration and convexity, they help you conceptualize what the risks are in your portfolio. For that reason they are often a good way to present results to senior management, especially if senior management does not have a quantitative background.

Chart 8 shows two price curves, one with the bailout and one without. If you look above the zero interest rate shift, you see that there is some additional value to the SPDA or additional cost, if you will, of having that bailout feature. The zero interest rate shift is a fully option-adjusted price starting with the current yield curve so that it encompasses interest rate movement going forward, and that's where you have the extra cost to the bailout. As you go to the right and left of the central point, you can see what happens to the cost of the bailout as interest rates move. The left side illustrates what happens when interest rates drop and the right side shows what happens when interest rates increase. When interest rates drop, you can see the gap between the SPDA with and the SPDA without the bailout getting bigger and bigger. They both go up in value, but the SPDA with the bailout costs much more. This is natural because, as interest rates drop, you're getting closer and closer to triggering that bailout. You have more paths and more of a probability that the bailout will get triggered so it's going to cost you more to have had that feature in your product.

CHART 8
SPDA PRICE CURVES

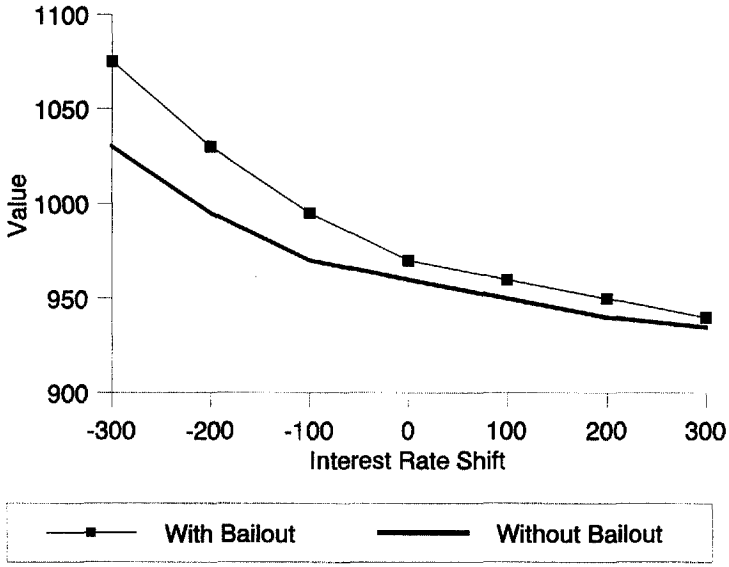

One of the nice things about the price behavior curves is that the slope and the curvature of the lines give you information. The slope tells you the duration and the curvature tells you the convexity. Here you can see, above the zero interest rate shift, which is our current yield curve, the SPDA with the bailout has a little bit higher slope than without the bailout, so it will have a slightly higher duration. The real trick comes in the convexity and that's demonstrated by the curvature. You can see how with the bailout, the product has much higher convexity than without.

Chart 9 applies the technique to both assets and liabilities. This could be a whole in-force portfolio or a model office projection for new business that projects the liabilities and the assets that will be used to support them.

The difference between the asset and liability lines is economic surplus. My asset line is my upper line; liabilities are below. I'd like to maintain economic surplus. If we start from the zero interest rate shift and move left, you can see what happens when interest rates go down. Economic surplus starts to increase a little bit, but then it decreases. The real trouble comes when interest rates go up and that's where you can see my lines crossing. I'm actually in a negative economic surplus position.

Once again, we look to the slopes and the convexities to get information about our asset/liability risk. If you look above the zero interest rate shift, we can see how my asset slope is much higher than my liability slope. The asset duration here is about 4.3 and the liability is about 1.1. The curvature is unbalanced as well. On the liabilities, we have positive convexity, positive curvature. On the assets, we have negative convexity, negative curvature. This becomes a great way to show senior management why convexity and duration mismatches are a problem. You can see how the duration and the convexity differences make these lines cross so that economic value is lost.

CHART 9
PRICE CURVES

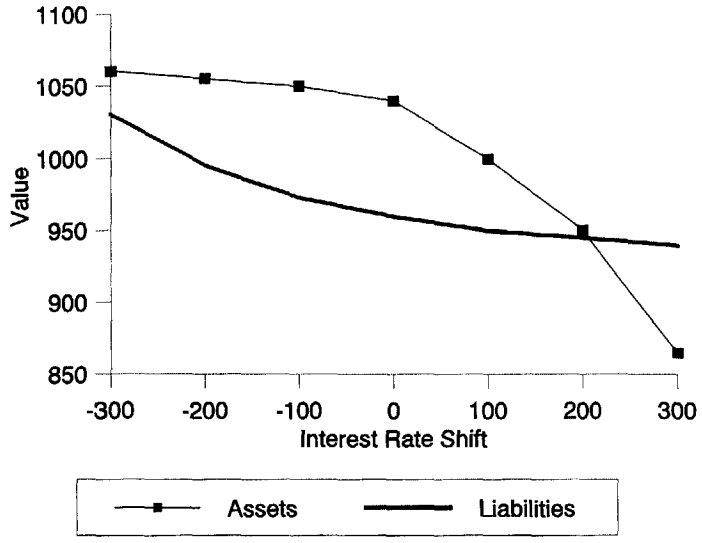

Price curves also provide you a good method for "fixing" your asset/liability risk that's changing the risk profile. The general idea is to line up these curves and make them "parallel." There are a number of ways to do this, just as there are a number of ways to fix the problem in real life. One way is to purchase asset hedges to change the shape of the asset curve to match the liabilities. We could also rebalance our asset portfolio to purchase a different mix of assets to accomplish the same thing. We can change our liability features as I showed you with the bailout. Changing some of the features in a product can change the shape of the liability curve to get it closer to match the asset curve. Finally, we can change our management strategy such as crediting rate strategy for the SPDA. In general, following the market tends to lower duration whereas portfolio type of strategies tend to increase duration. That's a very general comment because it really comes down to policyholder behavior to determine whether those durations and convexities go up and down with any strategy.

Chart 10 shows two hedges, an interest rate cap and a fixed-for-floating rate swap. The swap is used to correct the duration of my assets. When added to the asset portfolio, the slope of the assets is altered and gives us a better duration match with the liabilities. It's relatively easy because the price behavior curves are additive. Similarly, the cap is used to correct for convexity. It changes the curvature of the asset curve. Chart 11 shows that when I add these hedges to my asset portfolio in large enough quantities, I can actually line up my asset/liability price curves and preserve my economic surplus. This seems very easy. All we need to do is line up the curves and we're done. Where this falls short is in calculating the cost of those hedges. What is that doing to our earnings? I'm buying an interest rate cap and a swap and my current portfolio rates are going down because of that. Under a static environment, that cap is out of the money, so it's not even paying anything. I just pay a large amount of money for an interest rate cap and I'm getting no return unless interest rates move in the correct direction and with enough degree.

CHART 10
SPDA PRICE CURVES

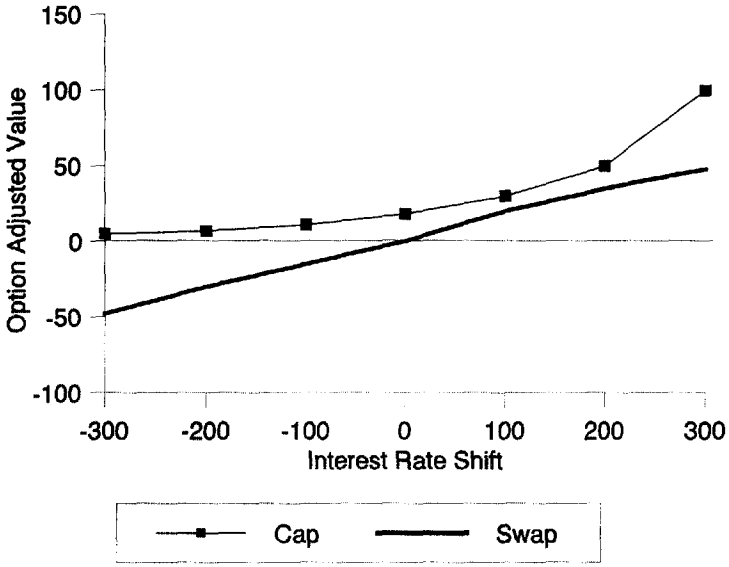

CHART 11
PRICE CURVES

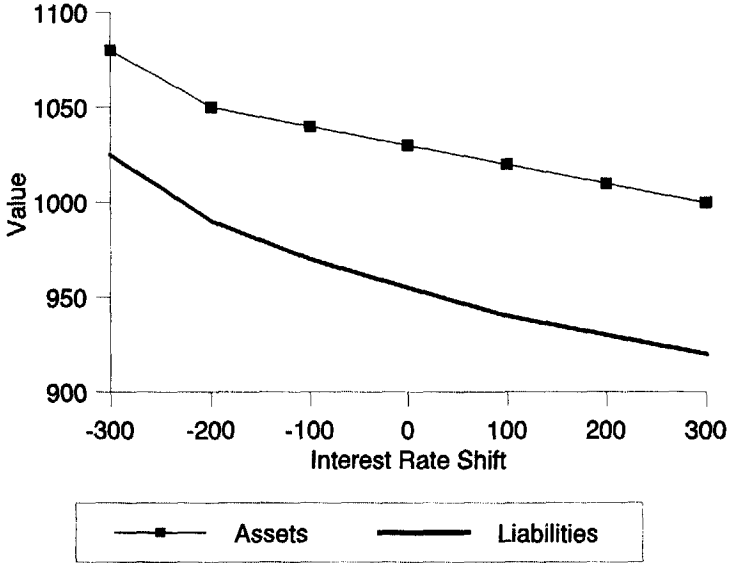

To find the earnings effects, you need to do the scenario analysis that Cheryl talked about. The general approach that I tend to use is the economic analysis which tells me what the risks are. I quantify those risks and make "corrections," but then, as a final check I do the earnings analysis to see how it affects my financials.

You can also analyze partial hedges. In Chart 11 I show a full hedge. You also want to look at partial hedges, for example, purchasing the cap and the swap in smaller amounts, so that these curves are not quite "parallel" but are better aligned than when we started.

Price curves are a good technique for your duration and your convexity risk, your "parallel" risk, but what about nonparallel shifts of the yield curve? For that we can use key rate durations. Key rate durations are a vector of numbers where each number evaluates the price sensitivity to a key rate shift. Each key rate shift is determined by breaking down a parallel shift into segments along the yield curve.

Chart 12 shows an initial curve, which is the bottom line and a parallel shift of 100 basis points above. That would be the curve used to determine duration and convexity for price-curve analysis. For key rate shifts, we take that increase of 100 basis points and break it down. For the first shift, we start out at the bottom yield curve and go up 1% at the short end of the yield curve. Then we drop down to the initial curve and continue to follow it. The second shift starts at the original curve, shoots up 100 basis points at our second point on the yield curve, comes back down at our third shift line, and continues to follow the lower curve. You keep doing this along the curve until you have it all pieced out and you have shifts that will show you the sensitivity of values to different points along the curve.

CHART 12
KEY RATE SHIFTS

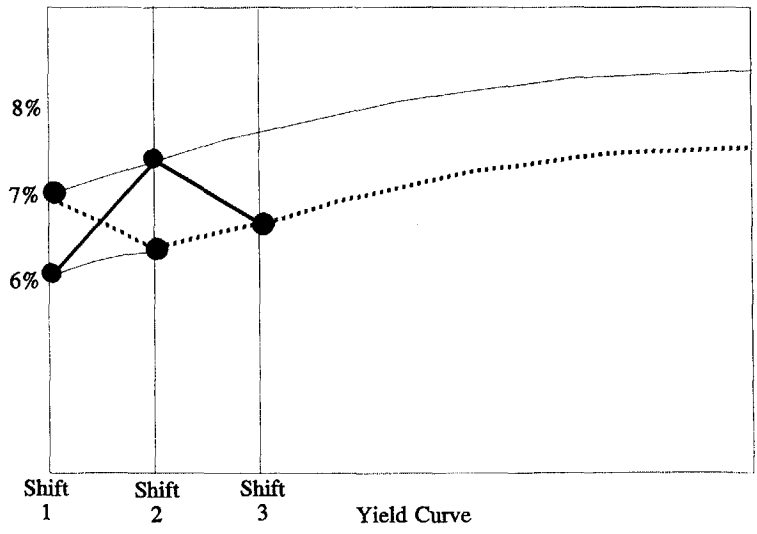

The key rates are defined so that the sum of the key rate shifts equals the parallel shift, and therefore, the sum of the key rate durations equals an effective duration. When I do this in practice, I typically use anywhere from five to eleven key rate shifts. Investment people tend to use 11 in practice. I think that's a bit of overkill. Five or six usually does the job for pricing purposes.

Table 3 shows key rate durations for my SPDA. I have an effective duration of 1.84 and it shows how that duration breaks down to the different key rate shifts along the yield curve. When I bring my yield curve up at the first shift, my SPDA value only changes by seven basis points. Add up the key rate durations and you get the effective duration.

TABLE 3
KEY RATE DURATION ANALYSIS

Shift	SPDA Duration
Effective	1.84
1 year	0.07
3 years	0.29
5 years	0.83
7 years	0.97
10 years	−0.43
20 years	0.11

Here you've got rather large key rate durations at the five- and seven-year points. This is due to large cash flows at these points. This is when my surrender charge is ending so I'm getting large lapses. The lapse is also kind of static. It happens under up and down scenarios. It's due to the surrender charge, rather than interest rate changes. Those all tend to increase duration. When we discount that large cash flow back at the five-year and seven-year term, it's going to make the value of my SPDA change quite a bit.

The other interesting number is the ten-year key rate duration. For a ten-year, we have a −0.43 key rate duration. This is due to increased interest sensitivity at this point on the curve. I have much more interest rate sensitivity because that surrender charge period has worn off. People are going to move. If I don't follow the market with my credit rates, they're going to move that money.

If I'm naive, I might create an asset portfolio to match my effective duration and it might look like Table 4. I might use a semibarbell approach where I have effective duration matched so I think I'm hedged. I match my duration and my convexity, but the key rate durations break down as follows. Because I'm using a semibarbell approach, the one-year key-rate duration is 0.67, the ten-year is 0.42 and there are smaller numbers in between, but this is where I have my large cash flows on my asset side.

TABLE 4
KEY RATE DURATION ANALYSIS

Shift	Portfolio Duration
Effective	1.84
1 year	0.67
3 years	0.18
5 years	0.14
7 years	0.22
10 years	0.42
20 years	0.21

Table 5 compares this against my liabilities, so you can see where the problems exist. I'm fine for effective duration. I have a perfect match but for key rate durations, I have large mismatches at the one-year, the five- and the seven-year and the ten-year rate. In this case, if interest rates shift and parallel, I'm okay. If there's a parallel shift in interest rates—if they both go up 100 basis points—then both the value of my SPDA and the value of my portfolio drop by 1.84%. So, I'm okay. But if there's a steepness shift in yield curve, a tilt or a twist, that's where the key rate durations can show you that you might be taking some risk. For example, if there's a 1% increase at the low end of the yield curve and a 25-basis-point increase at the high end of the yield curve, my liabilities would drop by 58 basis points, my assets would drop by 101 basis points and my economic surplus would actually drop by 43 basis points even though, from a duration and convexity effective duration standpoint, I'm matched.

TABLE 5
KEY RATE DURATION ANALYSIS

Shift	SPDA	Portfolio	Mismatch
Effective	1.84	1.84	0.00
1 year	0.07	0.67	0.60
3 years	0.29	0.18	−0.11
5 years	0.83	0.14	−0.69
7 years	0.97	0.22	−0.75
10 years	−0.43	0.42	0.85
20 years	0.11	0.21	0.10

Now if I match my key rate durations, my effective durations, my convexities, my price curves, and I feel like I'm in good shape from an economic standpoint, am I finished? Well, the answer is no because you still don't know what the financial impact is of those matches. You're still not sure what that's doing to your earnings. Furthermore, even if you hedged cash, you need to watch out for what the effect is on earnings. Hedging cash can often have the opposite effect upon your earnings due to the biases in our accounting systems. Again, for this you should perform the earnings analysis that Cheryl spoke of.

Finally, I'm going to show you one more technique to combine the option-pricing analysis that I'm showing you with the scenario analysis that Cheryl showed and that's called the option-adjusted value of distributable earnings (OAVDE). This was a technique developed by Dave Becker and documented in *Product Development News* in November 1991. We

take our option-pricing model and we apply it to the free surplus coming off of a financial projection. We incorporate accounting, benchmark surplus, increases in reserves that are needed, and our option-pricing model so that we look at up and down scenarios. We discount the free cash flows, the amounts coming off after all the reserves and surplus that we need are provided for. We probability weight those results across all the interest rate scenarios, discount them and we end up with an option-adjusted value of our distributable earnings stream. This is a very complex calculation. In practice, I've found it takes a lot of work to get your model to the point where this calculation can be credibly performed.

Chart 13 shows the option-adjusted value of distributable earnings with the hedge and without the hedge. It shows the earnings effect of having those hedges in place. Under a zero interest rate shift, I'm better off not having the hedge because the hedge costs me money under a static scenario. But as interest rates move up and down, I'm much better off having the hedge and taking a hedged, lower amount of earnings rather than an unhedged, higher amount of "static" earnings. The approach here is to test different types of "fixes" for your asset/liability risk, look at the effects and be able to assess what the risk return trade-off is.

CHART 13
OPTION-ADJUSTED VALUE OF DISTRIBUTABLE EARNINGS

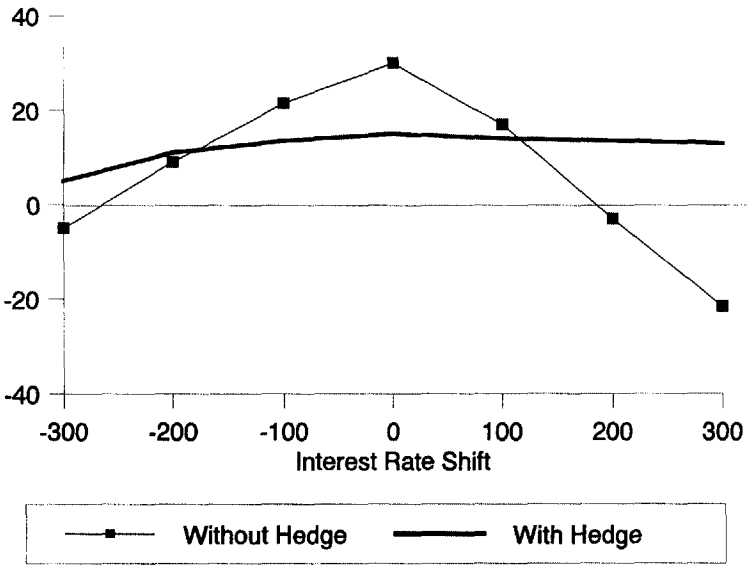

You can do this with percentage returns as well as absolute dollar amounts to find the option-adjusted return of distributable earnings. Without the hedge, I may be getting a 15% return, but with the hedge, I'm only getting 10%, but I'm locked in reasonably well with that 10% as opposed to taking a lot of risk with that 15% if interest rates move up or down.

There are some limitations to these models that we always need to consider. They do not manage our business for us. They really just give us insight into these risks. The current state of the art with corporate modeling doesn't allow you to take into account much of the factors that will happen in real life. Most models today, for example, can't account well for future management action. What happens when interest rates go along an undesirable path? Our models tend to say we're locked into a set of strategies and assumptions that we may not follow in reality.

Management is not ignorant. We must keep in mind that our models have these limitations. The other limitation is all the assumptions that we put in. Our models are only as good as our assumptions and there are a number of assumptions that we may not be completely comfortable with. First and foremost is probably policyholder behavior. That assumption is something that I know most people struggle with. We have very little data in terms of experience. Most people are taking a shot in the dark and using a best-guess assumption, which is all we can do. I think it becomes important to do sensitivity testing on these assumptions. We can't just use our models, do our analysis and say that we're done without looking at the effects of the assumptions that we're not certain of.

Chart 14 shows the effect of interest-rate-sensitive behavior. I show my original hedge and then I show what happens if we increase the sensitivity of our policyholders. Even though I purchased those hedges, I have more interest rate risk than I thought I did because my policyholders are more sensitive than I thought they were.

CHART 14
OPTION-ADJUSTED VALUE OF DISTRIBUTABLE EARNINGS

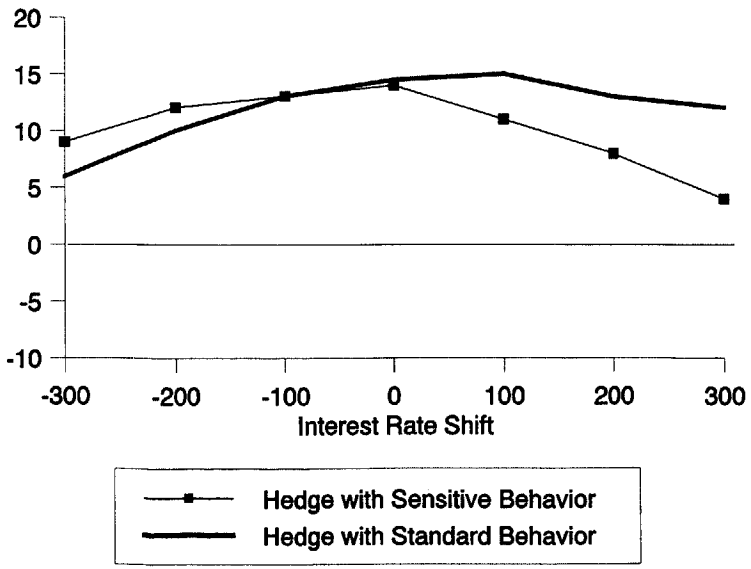

Let me now summarize. I recommend performing many different types of analysis to analyze your asset/liability risk. The OAS analysis explicitly values your options. The

price curves give you duration and convexity analysis. Key rate durations give you yield curve risk. The scenario analysis shows you the earnings effect of these risks. Option-adjusted value of distributable earnings combines the scenario school of thought with the option pricing. Finally, sensitivity testing is very important. You definitely want to do that.

The best way to do all of this is with the integrated system. You can have one set of paths and one set of analyses that not only do the option-pricing techniques, but also do the scenario techniques. The scenario analysis that Cheryl showed you can be done using the same paths and scenarios that are used to develop your key rate durations, your effective durations and convexities so that with one set of analysis, you can accomplish many different types of analysis. The integrated system is practical, it's flexible, and it gives you the best analytics. It gives you the least amount of maintenance and it helps improve run time.

MR. ROBERT A. NELSON: My question has to do with the interest rate models that are used. You talked about doing the option pricing and doing the efficient frontier. I will relate some of my own difficulties in implementing those techniques and get your opinions. One problem that I have is for the option pricing; we use arbitrage-free curves. That way I can rely on the cost of the hedge or the cost of options being accurate. When I go to do the efficient frontier, however, if I use arbitrage-free curves, one of the things you can quickly conclude is, there's never an advantage to buying anything other than the three-month bond. Why? Because according to an arbitrage-free model, there can't be. The ten-year bond is just the combination of holding all successive forward rates of three-month bonds. If you use the exact same curves, you always find that the maximum return with the least risk is to hold the three-month bond since you've gained nothing except a disintermediation risk by holding the longer bond. I'd like to get both of your views on how I get out of this conundrum. If it's pricing for options, I need arbitrage-free. To do an efficient frontier, it leads me to a rather strange conclusion.

MS. KRUEGER: We typically would not use arbitrage-free scenarios for efficient frontier analysis for that very reason. The technique relies on modeling future interest rate movement. The arbitrage-free scenario takes into account the current yield curve, but doesn't take into account future changes in the yield curve. Under efficient frontier analysis, you need to take into account all the changes in the future interest rates. I believe you need two different sets of scenarios to run the two types of analysis.

MR. GEORGE: When I perform the scenario-based analysis, I tend to look at explicit paths and explicit scenario conditions, so that I can, through maybe only 25 or 50 scenarios, be able to see the effects of many different kinds of movements in interest rates. I can do this because the paths are designed to give you reasonably optimal coverage of future interest rates with a limited number of paths. Sometimes I add a handful of deterministic scenarios to the paths because the paths are limited by whatever option pricing model you are using. In any case, these can replace 200, 500, or 1,000 stochastically generated scenarios that would be required to give you the same sort of coverage of future interest rate movement. This makes the analysis much more practical to implement. In addition, because I only need a limited number of paths, I can review the explicit earnings patterns over each path—something that is difficult to do in practice with stochastic scenarios, but something that I think is important to asset/liability analysis.

In any case, I don't particularly advocate performing efficient frontier analysis with this approach although I would look at the risk and return trade-off of any given strategy. I use option pricing to analyze the risk profile and suggest alternative courses of action for altering it like hedging, portfolio rebalancing, or changing product features or ongoing product management. Earnings analysis is used to determine the earnings effects of the alternatives, specifically the cost to current yield and the cost and bias of our accounting mechanisms.

The underlying problem that you are pointing out is that OAS models are based on the "expectations hypothesis" for future interest rates, that the current yield curve is an unbiased predictor of future rates. Historically, the yield curve has been shown to be a lousy predictor of interest rates. This is a concern for all OAS models. For this reason I perform analysis under a predicted interest rate curve as compared to the rates implied by the current curve at any one time. Option values and payoffs can be compared and interest rate bets can be taken.

MR. RICHARD COLE PRETTY: My question is for Cheryl, with regards to the spread adjustment approach of scenario testing. You mentioned four steps where you do your static pricing, solve for your spread there, and then you run it through your stochastically generated scenarios and adjust your spread in the third step. We've done that at our company, but we've never gone on to the fourth step and I'm not sure that I understand what the purpose of that fourth step is in going back to the static pricing. If you could elaborate on that, I'd appreciate it.

MS. KRUEGER: The assumption is you're trying to meet a certain ROI target and you have a product design that doesn't meet that target. Rather than change the spread, maybe you prefer to change other product features. So rather than making the product feature adjustments in your stochastic scenarios, you can just go back to your level scenario and make the product changes there, by pricing in the cost of the options that you've left in rather than going through the stochastic analysis. You wouldn't have to go back if you just want to make the spread adjustment.

MR. OWEN A. REED: I'd like to pick up on the question that was put to you, Cheryl. You never did really answer the question about how you generated the interest rates.

MS. KRUEGER: How we actually generate them?

FROM THE FLOOR: Who decides what the viewpoint should be? If they're not going to be market, then someone must be saying, I don't think the market is right. Who gets to decide that?

MS. KRUEGER: The decision is made based on your observation that the results aren't going to make much sense using the other technique. In terms of the follow-up question, are you asking, technically, how are they generated?

MR. REED: Yes, that's right. An efficient frontier series is really like a tall building on a balsa wood foundation. There's someone's arbitrary views on volatility and covariance and all this stuff and you've got another variable in there which is your future interest scenarios. The question is, how are they generated?

MS. KRUEGER: In the past, we have looked at volatility in terms of a 90-day and a 10-year rate and come up with some average volatilities, but certainly the interest rate generator is another assumption that goes into this type of analysis. The scenario generator that we use will generate both types of scenarios so if we want to do option-adjusted spread analysis, we make an adjustment to make the scenarios arbitrage-free. Other than that adjustment, the assumptions that go into creating two multiple sets of scenarios are the same. You have to make a volatility assumption and you have to start with an initial yield curve and then the only adjustment is to add an additional condition that requires that those scenarios be arbitrage-free. It's the same generator used in a final adjustment to make the scenarios arbitrage-free.

MR. REED: I guess I'll try one more time then. If the arbitrage-free model is a lognormal model, are you saying that you're generating interest rate scenarios for your cash-flow testing using a lognormal model as well?

MS. KRUEGER: Yes.