RECORD OF SOCIETY OF ACTUARIES
1995 VOL. 21 NO. 4B

STATE-OF-THE-ART RISK MANAGEMENT SYSTEM AND APPLICATION

Moderator: DOUGLAS A. GEORGE
Panelist: BRIAN C. TRUST
Recorder: DOUGLAS A. GEORGE

The teaching session will discuss constructing and using a comprehensive, credible risk management system with a concentration on C-3 risk. Applications will focus on practical analyses to assess risk and possible courses of action for rebalancing a company's risk profile.

MR. DOUGLAS A. GEORGE: I'm a partner with Avon Consulting Group in Avon, CT. I specialize in the financial analysis of life insurance companies for asset/ liability management and cash-flow testing, mergers and acquisition analysis, product development, capital planning, and other strategic analyses. In the asset/liability management area I have experience with both types of analysis that we're going to talk about today: scenario-based earnings analysis and the option-adjusted cash analysis. As a consultant, I've developed and implemented a number of risk management systems for life insurance companies.

Brian Trust is a vice president and investment actuary with Falcon Asset Management in Baltimore, MD. Falcon Asset Management is a wholly-owned investment advisory subsidiary of USF&G. Falcon manages portfolios and provides asset/liability management services for USF&G and external client insurance companies. Previously, Brian was with Fidelity & Guaranty Life Insurance Company, which is USF&G's life insurance subsidiary, where he did product development and asset/liability modeling.

Our discussion will proceed as follows: first we will talk about development and implementation of a comprehensive risk management system. This is the "how-to" portion of our session—how to construct and implement the system. Second, we'll discuss applications. We'll focus on asset/liability management, of course, and talk about two general approaches. We'll address option-adjusted analysis, which is a cash-based approach, looking at measures such as duration and convexity. We'll also talk about earnings-based analysis through the mean-variance approach; that is, the efficient frontier approach. I'll speak first and address the development and implementation of the system and the option-adjusted analysis and Brian will address the earnings-based analysis.

When I look to construct my risk management system, I first look to my actuarial software and I find that I have a number of concerns. The interest rate process seems to be rather simple. In general, the processes that are included with the actuarial software packages aren't as sophisticated as some of the ones available on Wall Street. They might capture all the nuances in my options. They might not price both short and long options accurately. Another concern is asset modeling. Creation and maintenance of an asset model in actuarial software can be difficult. Collateralized mortgage obligation (CMO) modeling can be limited, especially when it comes down to some of the trickier tranches.

Prepayment modeling can be less sophisticated. It might not take into account all the different factors that affect prepayment, such as age, seasoning, and seasonality.

Furthermore, the software often does not have a database so you have no automated way to get current collateral information and recent prepayment history. I need to go back and reconstruct my asset model each time I want to update my whole modeling system. Computer run time can be an issue. Actuarial software is often written in APL, so it tends to be slower. Consistency can also be an issue—consistency with the results that my investment people are getting. For example, if I'm performing asset/liability management (ALM) with my actuarial system, and it tells me my asset portfolio has a duration of 3.7, and the investment people claim that it's 3.1, and that's what their system is telling them, then I have a problem. It's hard to reconcile the numbers coming out of the two different systems. I might end up with a political battle on my hands. There can be credibility issues when trying to force investment people to manage their portfolios based on the results of actuarial software.

Next, I look to my asset modeling software—the trading systems that my investment people are using. The concerns here include liability modeling. It's tough to model liability products through asset forms. It's tough to get all the product features in there. I might not pick up all the underlying demographics of my liability in-force block. Accounting can be tough because the asset modeling systems don't usually come with statutory, GAAP, and tax support. We also don't have reserving. We don't have the income models. Asset valuation reserve (AVR), interest maintenance reserve (IMR), and risk-based capital (RBC) calculations usually wouldn't be present. Asset/liability interaction can be tough. It's tough to perform crediting strategies, especially complex ones, such as portfolio-based or investment-generation-based strategies. Future trading, such as reinvesting and selling and time horizon analysis, is usually not supported. If I want to incorporate new business into my model to do business planning, it can be difficult in asset systems.

The technique that I tend to use is systems integration. Here we're looking to obtain the best of both systems and merge them together. I start with my trading system used by my investment staff, but sometimes I use commercial software. In my personal experience, I've used software from Global Advanced Technology (GAT) and Capital Management Sciences (CMS). I believe that other commercial systems are out there, and you also can use internal systems, homegrown systems that are developed for your company that your investment people use to project cash flows on assets.

On the actuarial side, I have the same situation. I use PTS to do this. I know you can also do this with TAS, and I believe other actuarial systems would support this process. You can also use homegrown, internally developed actuarial systems.

I start with my asset modeling system (Chart 1), usually a Wall Street system, when I need to incorporate my interest rate process. If I'm doing option pricing analysis, I'll need an option pricing model. I'll go to my Wall Street model because that system tends to have a more robust option pricing model; it gives me better pricing for both short and long options. If I'm doing earnings-based analysis or stochastic testing, I would need stochastic scenarios and a stochastic generator. You can take that from a variety of places. The interest rate process I'm using (which, of course, would be based on the type of analysis I want to do), will be connected to my asset modeling system. Here I'll pick up my asset database. I'll pick up my more sophisticated prepayment modeling from the commercial system from Wall Street. I'll get the deal structures that I need for my CMOs.

CHART 1
SYSTEMS INTEGRATION

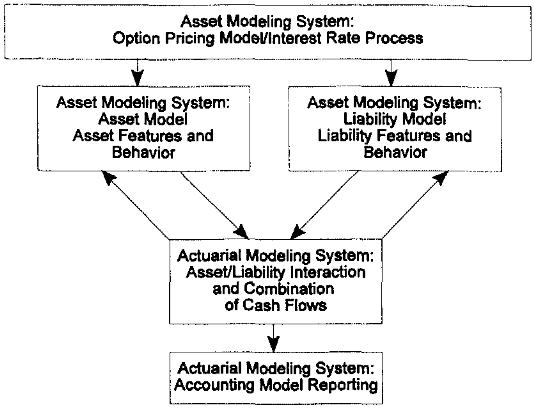

On the liability side, I'll hook up the interest generator to my product model in my actuarial software. Here I'll pick up exact product modeling, in addition to the liability behavior, demographics, and my underlying in-force block.

My experience with this is that this diagram tends to get bigger. I tend to have many systems hooked up on the asset side and on the liability side. You might have commercial systems on the asset side, as well as some homegrown systems. Commercial systems might be used to run CMOs and mortgage-backed securities. You might have a homegrown system to run derivatives or commercial mortgages. On the liability side, you might have commercial systems for single-premium deferred annuities (SPDAs), universal life (UL), and even traditional life products. But they might not have good support for GICs or terminal funding. There you might have your own internal homegrown system that you'll also want to hook up to the overall integrated system.

Next, our integrated system is going to need capability for asset/liability interaction. Here I usually use the actuarial system, especially if I want to do earnings-based analysis. The actuarial system will support the future cash-flow interaction that I need. It'll do the time horizon analysis that I need. It'll support the buying and selling of assets in the future. It'll support the more complex strategies, such as a portfolio-based crediting strategy or the investment generation crediting strategy. I can also use the investment system for this purpose, if I'm limiting my analysis, to option pricing only. Here we don't need the accounting. We don't need the time horizon. All we'll be looking at is time zero analysis, and in this case it's often best to use the investment system to do the merging of the asset/liability cash flows.

At this point we often run into some modeling limitations. Current modeling systems have limitations, especially with modeling strategies, such as reinvestment and crediting strategies. Our current software tends to lock these in as we set up the model, and no matter what happens to interest rates or economic scenarios, we still are stuck with the same strategy, and that's not the way we run our companies. I tend to modify the systems at this point to try to better mimic reality, to reflect the fact that we do have management in place and that if interest rates go one way or the other, or other conditions change during the next few years, we will take management action other than that which our model might lock us.

Finally, we need accounting. We need an accounting system if we want to do our earnings analysis. Here we can pick up our statutory and GAAP support. We'll go to the actuarial system for this function. We also need a reporting methodology, and that can be either system, although it often flows well to just use the actuarial system to report results.

I have my integrated system along with some concerns. First, pieces of my overall model can be missing. For example, the asset modeling system typically doesn't come with an asset default model. There's no support for C-1 risk. Also, accounting for my assets can be limited. Second, there can be some consistency issues. Interest rates can be different—the definition of interest rates in the different systems. One might be based on a bond equivalent yield, another on annual effective, or one might be spot, while another is forward. Cash-flow timing can also be an issue. One system might have quarterly cash flows while another might be monthly and still another might have exact-date logic.

My experience has been that each of these issues can be resolved adequately. The missing models can be added, and usually rather easily. The C-1 models that we have in our actuarial software are rather simple anyway, and just adding them to our asset cash flows is not very difficult. Accounting can also be dealt with. Oftentimes the investment systems do have accounting support, it's just that people aren't aware of it. Our investment people are running these systems, and they're not interested in the accounting. If the support is not there, we do need to do some approximation, and my experience has been that we can get adequate results by using an approximate method, but usually there will be good enough support in the investment system. It just might be a matter of loading up book values to give us the earnings information that we need to support earnings analysis.

The consistency issue is rather straightforward. Interest rates can be converted if there's a definition difference in one system versus another. If the cash-flow timings are a little different, cash flows can be brought forward so that they are consistent from one system to the next. Interest rates can be converted rather easily so that we do have consistency.

Once we've resolved our concerns, I see a number of advantages to the integrated system. First and foremost, we're getting the best analytics possible for both sides of the balance sheet. Our asset software is modeling our assets the way it should. Our liability software is modeling our liabilities. We're getting the best analytics available.

Second, we're getting the best interest rate process available, and we're getting a consistent one because the same process is feeding both our assets and our liabilities. If we want to calculate durations or convexities, we're going to get an apples-to-apples comparison between our asset and liability durations. If our assets are duration 4.5, and our liabilities are 3.5, then we can be sure that the mismatch is 1.0. This is opposed to a situation in which you have different systems and different interest rate models are producing a duration for assets versus a duration for liabilities and we're going to model bias so we can't be sure that our mismatch is 1.0.

Third, model creation and maintenance is as easy as possible. Certainly there's a great deal of legwork in creating our corporate models, but because our asset modeling software is modeling our assets and our liability modeling software is modeling our liabilities, we're making it as easy as it can be. The systems are flexible, and I think we'll show you that with some of the analysis that Brian and I will talk about. They're also consistent. There's going to be consistency with the independent analysis that our investment people are doing and with what our actuaries are doing. We're not going to run into those political battles. If we try to implement a change to our investment strategy due to the results of our risk management system, our investment people will buy into it because we're using their system to project the assets. We're not going to have the political battle where we're trying to tell them to change their investment strategy based on the results of actuarial software. Finally, computer run time can be improved because the systems tend to run simultaneously. We're not lining up everything in a queue, with our investments and liabilities all running one after the other. We can get some simultaneous running of results, and that can improve run time quite a bit.

There are a number of applications for a risk management system. We can use it for capital planning, where we want to look for the optimal capital structure for our company. We can analyze RBC. We want to look at what capital is needed to maintain our ratings, what kind of return we need to get, and that balance of capitalization versus our return. We can add new business to the model and do financial forecasting; that goes hand-in-hand with the capital-planning function. We can do compensation analysis, looking at different compensation structures for our field force. Pricing can be done with our integrated system, and we can incorporate asset/liability analysis into product development. Our system can be applied to the macropricing algorithm that seems to be popular these days. If you're looking for the optimal production profitability trade-off, I believe that asset/liability analysis should be incorporated as part of that.

Performance assessment can be done with our integrated system. We can support the total return approach and value-added analysis—two types of analysis that many people are advocating these days.

Finally, there's merger and acquisition appraisal. I could argue that this is the best way to do an acquisition job. Not only are we getting good analytics, but the model creation is as easy as possible, and that can be especially important for acquisition analysis because of the deadlines that are critical in an acquisition job. You have one week to appraise a company. You'll want to get the system up and running as quickly as possible, and this is one of the best ways to do that.

Of course, today we're going to talk about asset/liability management, and the integrated system can support a number of different approaches to asset/liability management. We're going to talk about two general approaches. First, I'm going to talk about option-adjusted spread (OAS) analysis, and then Brian will talk about earnings-based analysis, otherwise known as the efficient frontier approach.

Starting with the OAS, I want to show you a demonstration on an SPDA (Table 1). To calculate an OAS, we're going to use the all-in-cost cash flows. This includes the benefits, expenses, premiums, and commissions. It doesn't include taxes because this is a cash-based analysis, and to do a tax calculation you need to have a tax module and an earnings module, and it does not have that. All the values will be pretax.

TABLE 1
OAS ANALYSIS–SPDA

Run	OAS*	Static Spread*	Option Cost*
Base	68	48	20
Increase initial rate 50 basis points	88	66	22
–Difference	20	18	2
Reduce surrender charge 1% per year	81	56	25
–Difference	13	8	5

*in basis points

The first thing we calculate is the static spread (middle column). Under our base run we project the cash flows, the all-in-cost cash flows of our SPDA along the forward rates underlying the current yield curve. We solve an equation so that the present value of cash flows equals the premium paid for the SPDA, and when we do that, we find in this case that a spread of 48 basis points above current Treasuries solves that equation. So the present value of cash flows calculated at a discount rate of 48 basis points above Treasuries equals the premium paid for the SPDA. That's our static spread because it's based on a static yield curve.

Second, we expand this to do our OAS calculation, and here we project along a full range of interest rates. We include not only the current yield curve but also paths that go up and down as well. Due to the interest rate sensitivities in our model, some of the cash flows under some of the paths will occur sooner under the OAS calculation than under the static yield curve calculation. This tends to increase the value of those cash flows. In this case, an OAS of 68 basis points above Treasuries is the rate needed to equate the present value of cash flows averaged across all interest rate scenarios to that premium.

What does this tell us? First, if I looked at my assets, I know I need to earn a 68-basis-point OAS, in order to break even, averaged across all interest rate scenarios. I need to earn a 68-basis-point OAS, averaged across all interest rate scenarios, just to achieve a zero profitability on my assets. Second, I calculate my option cost, which is the difference between my OAS and my static spread, and find it to be 20 basis points. This is the value of the option that I've granted my policyholder to surrender at book value. Because I'm giving him or her an option to surrender at book value, it's costing me 20 basis points. Third, I can test different asset/liability strategies. I can test them to maximize the difference between the OAS that I earn on my assets and that which I pay on my liability. I can test different combinations of investment strategies and crediting strategies so I can maximize that difference.

Finally, I can explicitly price the options in my liabilities and products, and this is arguably the best use for OAS analysis. I show an example of increasing the initial credited rate by 50 basis points, and in this case the OAS goes up by 20 basis points. I can also look at decreasing my surrender charge by 1% a year, and here my OAS goes up by 13 basis points. Notice how on the surrender charge in the far column of Table 1, our option cost has also increased by five basis points, and this is due to the fact that we're giving the policyholder the option to surrender at a higher book value. My option cost actually increases by five basis points due to that.

With this analysis I look at the profitability of the products in my portfolio. The marketing people might come to me and say, "Our SPDA is a real dog, we need to beef it up somehow. We're indifferent between increasing the initial rate by 50 basis points or reducing the surrender charge by 1%; that is, we think we'll get equally increasing productivity based on either one of those. Which one should we do?" My answer is to reduce the surrender charge 1% because I realize that's only costing me 13 basis points of OAS versus increasing the initial rate by 50 basis points, which costs me 20 basis points of OAS. I can also look at the different products in my portfolio by using the same technique and assessing what their credited rate should be. If I have a couple different SPDAs and they have different product features, let me equate the costs of those. Let me equate the OAS so that I can solve for the initial credited rate that each one should have. If one has a little bit different surrender charge, another has a different bailout feature, and so on, I can solve for what the credited rate should be on the different products, so that I don't care which one my field force sells. I'm going to get equal profitability based on either one, averaged across different interest rate scenarios.

I show one or two more examples in Table 2. I'm working with the bailout. My "base" product did include a bailout option for five years, and if I look at my static spread, when I remove the bailout, it doesn't cost me anything—I still have a 48-basis-point static spread. If you think about it, that's the way it should be because the bailout wasn't triggered under the static yield curve. My analysis will tell me that it doesn't cost me anything. Obviously, that's not reality, and that's why you do the OAS analysis. The bailout is triggered under certain interest rate scenarios. By removing the bailout, I see that the OAS goes from 68 basis points down to 58. Notice how the ten-basis-point difference transferred completely over to the option cost. There's a ten-basis-point option cost difference as well. This is conceptually true because the bailout is a pure option. It's only triggered when interest rates move, so it should be a pure option cost.

TABLE 2
OAS ANALYSIS—SPDA

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

*in basis points

I will show one more example in which I removed the bailout on an in-force block. In this case my bail-out has been in in-force already for three years, so only two more years were left in the bailout option, and it only cost me four basis points to remove it. You can see how OAS costs change as you get closer or further away from an option starting or ending. If you think about it, the SPDA with and without the bailout is the same product after five years, so there wouldn't be any option cost difference once you've passed the bailout period.

The OAS analysis is good because it shows you the option costs under current interest rates. It takes into account the fact that interest rates can go up or down, but it doesn't show you what happens if interest rates do go up or if interest rates do go down. For that we can use a technique called price behavior curves (Chart 2).

CHART 2
SPDA PRICE CURVES

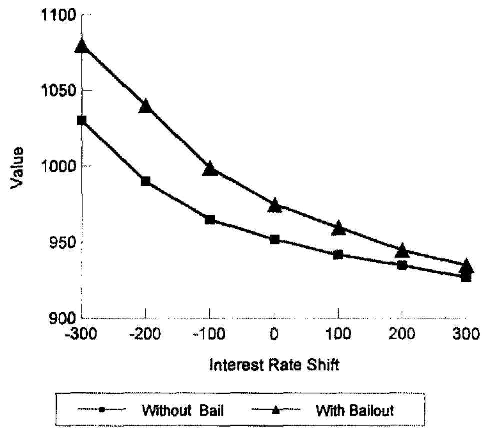

The price curves are good because they give you a picture of your duration and convexity risk, most of your asset/liability risk but not all of it. They can help you conceptualize duration and convexity risk so they can be a great tool for presenting results to senior management. They give you a picture of your risk profile. In my example, I show an SPDA with and without a bailout option. You can see how the bailout costs more as interest rates go down to the left. The bailout tends to be triggered more often, and the bail-out does cost more. Chart 2 illustrates our risk by looking at the slope and the curvature of the two lines. The duration of the lines is indicated by the slope. If we look above our middle point, our zero interest rate shift, you can see how the slope of the SPDA with the bailout is slightly higher than without the bailout, indicating that the duration be slightly higher for the SPDA with the bailout. Convexity is the same way because convexity is indicated by the curvature of the lines. Once again the SPDA with the bailout is more curved, has higher curvature and, therefore, higher convexity.

I can apply this to my complete assets and my liabilities once again because I have the integrated system. If I do that, I might get a picture that looks like Chart 3. The difference between the asset and liability curves is considered to be economic surplus, and you can see how that changes as interest rates go up and down. As they go up, we actually reach a point where we have negative economic surplus; not a very desirable position. Once again the slopes and curvatures of the lines can be a good way to present results to senior management and show them why duration and convexity mismatch can be a problem. In this case, my duration of my liabilities is 1.1, and the duration of my assets is 4.3. You can see by the slope above the zero interest rate shift how that duration mismatch is going to hurt me. The same is true with the curvature for convexity. The curvature in my liabilities is positive, and in my assets it's negative; that directly translates to positive and negative convexity. Once again it can show you and senior management why duration and convexity mismatch can be a problem.

CHART 3
SPDA PRICE CURVES

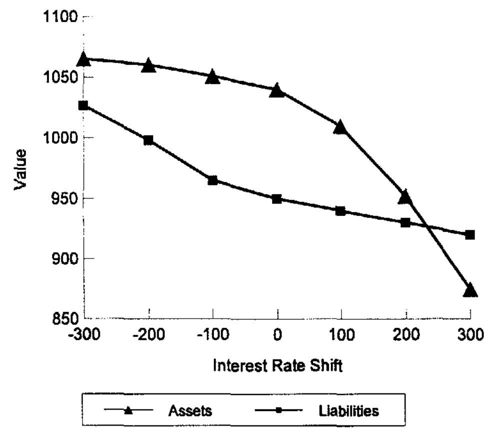

Another thing that price curves are good for is they provide you with a method to rebalance your asset/liability risk profile. They give you an approach. The general approach is to preserve economic surplus, and you can do that by trying to line up these curves. The curves aren't that hard to line up because they're additive. If I have a couple different securities in my portfolio, the total curve for the portfolio is the weighted value of the individual curves for the individual securities. I can make a number of different fixes to my asset/liability profile here. I can rebalance my asset portfolio. I can change my liability features at least on my new products going forward, and we saw how to do that with the bailout—how that changes the risk profile. I can change my in-force management strategy along with my editing strategy on my liabilities; that will change this picture. The fourth approach would be to do some hedges, buy some derivatives, and hedge my asset portfolio a little bit, and Chart 4 shows how you might do that.

In this case I bought a cap to help increase convexity and a swap to help change duration. You can see how the cap will help me increase the convexity of my assets because it has the curvature that I need to correct that asset line. Once again, the swap has a negative duration, so it will bring down the duration of my asset portfolio to help "correct" it.

I have shown a case of a full hedge in Chart 5. I bought enough of the cap and enough of the swap to completely preserve economic surplus under parallel interest rate movements. I'm not sure that this would be a desirable position, but I'll get to that later. Also, you might want to look at doing partial hedges if you're taking this type of an approach. Buy a small amount of both the cap and swap so that you can correct the lines but perhaps not completely line them up. The one limitation that these curves have is that they only work for parallel shifts in the yield curve. We're looking at duration and convexity risk here. I'm not sure what will happen if my yield curve moves in a nonparallel fashion. I'm not sure what the effect of yield curve risk is. Even if I line up these curves, it doesn't tell me what happens if the yield curve moves in a nonparallel way. For that we can use a technique called key rate durations (Chart 6).

CHART 4
SPDA PRICE CURVES

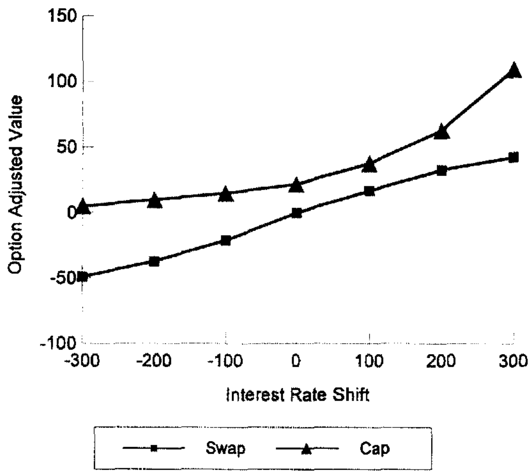

CHART 5
SPDA PRICE CURVES

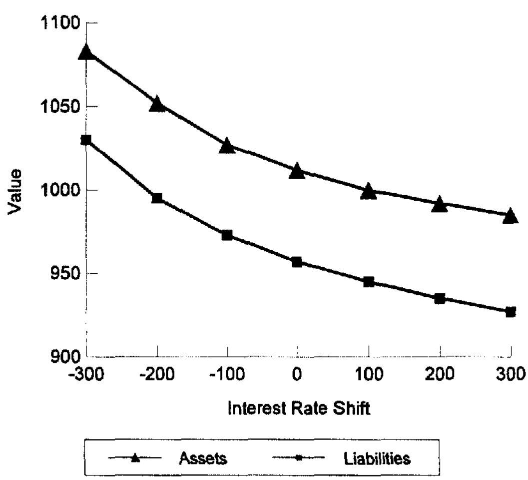

CHART 6
KEY RATE SHIFTS

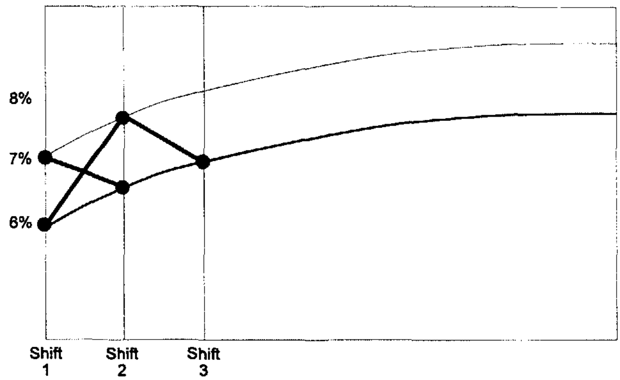

Key rate durations are a vector of numbers and each represents the price sensitivity to a key rate shift. Each key rate shift is found by breaking down a parallel shift into segments along the yield curve. I have an initial yield curve on the bottom and I have a parallel shift up top. The first two key rate shifts are shown. Shift 1 is an increase in yield for investments with short maturity. Shift 2 is an increase in yield for investments with a slightly longer maturity. In this case we'd have further shifts going down along the yield curve so that we would slice up that complete parallel shift into key rate shifts. We slice it up in such a way so that the sum of the key rate shifts equals the parallel shift and, therefore, the sum of my key rate durations equals my effective duration. I've seen people in practice calculate anywhere from 3 to 11 key rate durations based on 3 to 11 key rate shifts. I tend to use six or less. I think any more is close to overkill. You're not going to see that much independent interest rate movement.

Table 3 shows an effective duration of 1.84 on my SPDA, and we can see how that breaks down into the different key rate durations. A number of factors affect how the key rate durations come out. In this case you can see at the five- and seven-year points that we have relatively high key rate durations. This is due to large cash flows that occur at these points. These cash flows are based on the fact that the surrender charge expires, and are not based on interest rate sensitivity. They're going to happen no matter what happens to interest rates. This tends to produce higher key rate durations. At the ten-year point we might actually get a negative key rate duration, and that's slightly counterintuitive. Here the value of the SPDA actually increases if the ten-year interest term increases. This is due to two factors. There's a great deal of interest rate sensitivity at this point on the yield curve because we are past the surrender charge period. People tend to move their money. The other reason is when we increase interest rates at the ten-year point, it tends to increase the crediting rate at earlier points because our crediting strategy is based on the shape of the yield curve. Higher rates at the ten-year point tend to increase the credited rates that we apply to our SPDA, and that tends to increase the value of the SPDA and produce a negative key rate duration.

If I'm naive, I might invest in a portfolio with durations similar to those shown in Table 4, where I show where I've matched durations perfectly. I have an effective duration of 1.84 just like I need, so I'm in great shape for a parallel shift. But I haven't considered my key-rate durations. I've used sort of a semibarbell approach. You can see the high key-rate durations at the one-year and the ten-year point, but these don't match up well when I compare them with my SPDA key-rate durations (Table 5).

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

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
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

You can see how there's a match for the effective duration. We have a zero mismatch on the effective basis, but at different points on the yield curve we can have big mismatches— here at the one-, five-, seven-, and ten-year points. If there's a parallel shift in the yield curve, I'm still hedged in this situation. The value of my assets and liabilities either go up or down by the same amount under a parallel shift, but, if there's a nonparallel shift, that can be trouble. For instance, in this case, if I had a steepness increase in the yield curve and we increased interest rates 1% on the low end and 25 basis points on the high end, in this particular example our liabilities would drop in value by 58 basis points. My assets would drop in economic value by 101 basis points, so my economic surplus would actually drop by 43 basis points, even though I thought I hedged myself.

If I match my durations, convexities, key-rate durations, and price curves, does that mean I'm done? Does that mean I'm perfectly hedged and that I don't have to worry about this anymore? Not really, because we haven't considered the cost of doing the rebalancing. We really haven't looked at what I am paying for these hedges. Can I make money even if I hedge myself? The answer is typically no. Basically what happens is if you do a complete hedge, you hedge all your earnings away. Today's environment is very competitive. Our industry's very competitive right now, and you can end up locking yourself into a hedged profit, but the profit can be zero or negative, and you're going to earn it under all interest rate scenarios.

Another thing that can actually happen is by hedging cash you can't be sure what the effect will be to your earnings. There are biases in our accounting systems. Statutory and GAAP biases do exist, so that if you hedge cash, you might not be hedging earnings completely. Sometimes you can even have an opposite effect on earnings. For this reason you want to do earnings analysis in addition to this. Some people might say that this is all you need to do, but I think you need to do earnings analysis in addition to this type of analysis. You want to look at partial hedges or partial fixes as well as full fixes, about which I previously talked.

The way I tend to approach doing this type of analysis is to use the option pricing and the cash-based analysis to quantify my asset/liability risk, to get a feel for it, and also to design ways to fix it or rebalance it. The price curves and using the adaptivity of them can help you design ways to go about developing hedges or rebalancing or developing new crediting strategies or combinations of these different things that might work for you.

I tend to outline the approaches to rebalance the asset/liability profile through option pricing but then do the earnings analysis to assess the trade-off. This shows the risk return trade-off that you're getting by doing each of these different types of rebalances and "fixes."

Once again, the integrated system supports both of these and is going to be the best way to perform these analyses. At this point, Brian will talk about the earnings-based analysis, more specifically the mean-variance or efficient frontier approach that is widely used today. He will also talk about the importance of sensitivity testing for all of this analysis—not only the earnings but also the cash.

MR. BRIAN C. TRUST: As Doug described, I'm going to show you the approach followed by my firm, Falcon Asset Management. This process is broken up into five major steps. They are: modeling of cash flows, generation of economic scenarios, optimization, sensitivity testing, and portfolio selection. I will define each of these in more detail and give you an example by using a hypothetical SPDA case study.

In the first step, modeling of asset/liability cash flows, we look at the liability assumptions. For new business models, I start with best-guess estimates developed from the pricing assumptions provided by the pricing actuaries. If the analysis is for an in-force block, we use the best-guess estimates used in our forecasting and budgeting process as the basis for the liability assumptions. Because most of our forecasting and budgeting is done in a static interest rate environment, I have to add in policyholder behavior characteristics that reflect the dynamics of interest rates with lapses or premium payments. I also have to reflect company behavior; that is, how the company sets its credited rates. For asset modeling, I prefer to use our fixed-income analytics software package. We use CMS's Bond Edge. The portfolio managers are much better at modeling the assets because they do just that on a daily basis. As we purchase new assets, they make sure that the bonds are modeled appropriately in the systems that they are using to track their portfolios. I project the asset cash flows (including market values and book values), manipulate the data, and bring these cash flows into the actuarial asset/liability software. I also have to model the characteristics of the asset classes that I'm going to use for reinvestment purposes. My preference is to stick with the most generic asset class possible. First, I am uncomfortable with modeling some of the exotic asset classes that are brought to me. Frequently, the portfolio manager will ask me how we can possibly fit this new asset type into our investment strategy. The fact that it is an available asset today does not guarantee that it will be available next week, much less 10 or 20 years into the future. I prefer sticking with noncallable and callable corporates, residential mortgage pass-throughs, and some of the simpler CMO structures.

To model negative cash-flow strategies, I prefer to use a strategy that sells assets rather than to use a borrowing strategy. Therefore, I have to define a sale priority to the different asset classes. I have to take into account any trading constraints that we might have. For example, any of the assets that are in the *Financial Accounting Standard (FAS) 115* held-to-maturity category should not be sold in the asset/liability models. If I am modeling callable assets, then I will also need a prepayment model.

The third key area is the target surplus assumption, which needs to reflect an amount that is appropriate for the asset/liability risks that I am taking. For our SPDA case study I will assume that target surplus is equal to 200% of the NAIC's RBC formula.

Other case study assumptions include:

- Five percent surrender charge for the first five years
- Credited rate is portfolio yield less a 200-basis-spread target spread
- Two percent interest bonus in the first policy year
- Three percent minimum interest rate guarantee
- Base lapse assumption is 5% during the surrender charge period, 50% in the year after the surrender charge expires, and 10% thereafter
- Excess lapses follow an arctangent function capped at five times the base lapse assumption
- One percent of assets are kept in cash
- Positive asset cash flows are invested according to a fixed percentage mix
- Negative cash flows are covered by selling assets in the order of shortest to longest maturities
- All assets are available for sale under *FAS 115*

In Step 2 we generate the interest rate scenarios. We use 100–200 scenarios generated by an economic simulation model. Numerous input parameters must be calibrated when you use these models, including factors such as volatility and mean reversion. You should rely on any economic forecasts to some extent. Also, you need to consider historical data. To test your assumptions, you have to run your models to see how well the generated scenarios fit with what has actually happened through history. I want to emphasize that the scenarios I'm using here are realistic interest rate scenarios, not arbitrage-free scenarios used for option pricing and market valuation. I am interested here in analyzing risky strategies, therefore I need interest scenarios with as many realistic scenarios as possible, including yield curve inversions. Scenarios generated by many of the models do not have these characteristics.

Another key economic input is market yield spreads. I look to see what current market spreads are available, and I may use those as my initial conditions. But I also look to historical data to see where spreads might be moving in the future. Then I model reversion of spreads to those long-term averages. Also, I model market yield spreads as a function of the yield curve because spreads historically have tended to widen as interest rates go up and narrow as interest rates go down.

To model default loss rates for the different asset classes, I look to studies done by Moody's and Standard & Poor's (S&P). I use other studies for commercial mortgages. I will make an adjustment for actual company experience and expectations, especially with commercial mortgages and high yield bonds which are, of course, the higher default assets. The inflation assumption is related to the Treasury yield curves. For this case study, I have assumed that the inflation rate is the 90-day Treasury rate minus 150 basis points.

If I am modeling assets with any embedded options, I need some type of option pricing model for valuing interest-rate-contingent cash flows because market values are needed for assets to be sold. For the case study, the initial rates are equal to the ultimate mean reversion rate levels, which are 6% for the 90-day and 8% for the 10-year rate.

Step 3, optimization, is the really interesting part. We use what is called a mean-variance approach (Chart 7). When expressed graphically, the *y*-axis measures reward. We're looking for high reward. I use the median return across the 100 scenarios for that measure of reward. The *x*-axis shows risk. Here, we look at some measure of the variance of returns across the scenarios, such as the standard deviation. The mean variance approach was developed and championed in the 1950s by Harry Markowitz, who is considered to be the father of modern portfolio theory. His concepts related to portfolio selection state that the portfolio manager should be accountable for producing the highest return for a stated level of risk. The purpose of this mean-variance graph is to find portfolios that have either the highest return for a given level of risk or the lowest risk for a given level of return.

In Chart 7, Portfolio B would be considered an inefficient portfolio in this context because Portfolio A to its left has the same expected level of return but has less risk. Similarly, Portfolio C has the same level of risk but a higher expected return. The efficient frontier consists of all portfolios on the curve that go through Points A and C and have the best trade-off between risk and reward.

CHART 7
MEAN-VARIANCE APPROACH

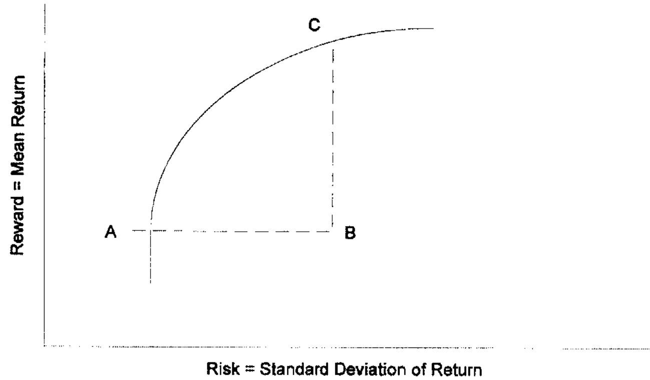

For optimization, sometimes it is hard for management to relate to a numeric measure of variance, so another common way we present the results is through a distribution graph showing the different percentiles for the reward measure. You can see in Chart 8 that Strategy C has a higher median return, represented by the horizontal center line. This strategy has a higher upside return at the 90%, but it also has a much more downside risk at the 10%. This graph is easier for management to understand and helps them to determine how much risk they are willing to accept. For the SPDA case study, I have defined reward as the median internal rate of return (IRR) across the 100 scenarios as calculated over a 20-year time horizon. I have assumed that the company's goal is a median return of 12%. For risk I used the standard deviation of the present value of distributable earnings at the 12% hurdle rate. We will also look at the percentile distribution bar charts and consider the 10% results as a measure of downside risk.

CHART 8
DISTRIBUTION OF RETURNS

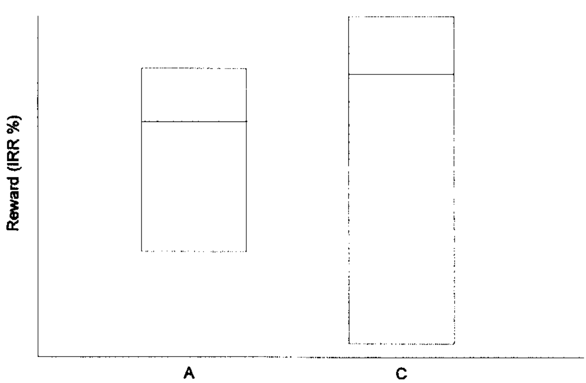

For this case study I ran more than 250 different investment strategies, combining four different asset classes (three-, five-, seven-, and ten-year investment-grade, noncallable, corporate bonds) in all combinations of 10% increments. You can see in Chart 9 that I have selected the efficient portfolios at various levels of risk and connected the points. These portfolios cover a wide range of maturity mixes. The net portfolio yield ranges from 7.38% for the least risky portfolio to 7.84% for the most risky portfolio. Once again, the crediting rate was calculated by subtracting the 200-basis-points target spread from the portfolio yield. Note that the initial portfolio durations range from 3.1 up to 5.7. All these portfolios are "efficient," but company management must decide how much risk they are willing to take before a portfolio can be selected.

CHART 9
EFFICIENT PORTFOLIOS AT VARIOUS LEVELS OF RISK

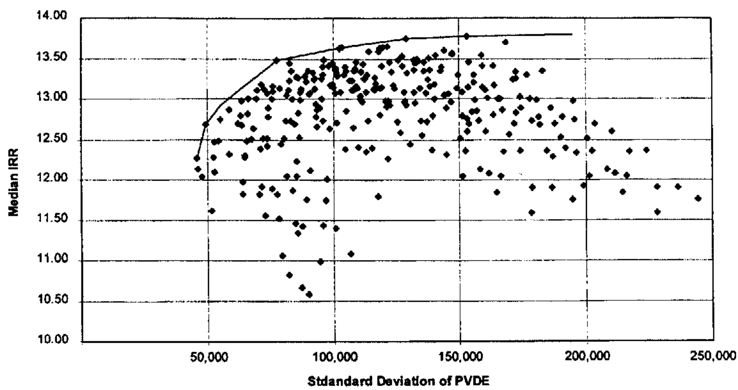

Allocation	7	11	89	173	196	240
3 year	70%	60%	20%	0%	10%	0%
5 year	30	40	40	30	30	30
7 year	0	0	40	70	30	30
10 year	0	0	0	0	30	40
Net yield:	7.38	7.42	7.65	7.71	7.77	7.84
FY credited rate:	7.38	7.42	7.65	7.71	7.77	7.84
Duration:	3.1	3.3	4.3	5.0	5.3	5.7

If I were making the decision as to which portfolio along the efficient frontier to choose, I would probably lean toward Strategy #89 because it is near the point where the line's slope is steepest. If you go out to Strategy #173, the medium IRR increases slightly (about 15 basis points), and the risk measure rises from about 75,000 to more than 100,000. So you're picking up almost 33% more risk for only 15 basis points of more return. This looks at the mean-variance approach only.

On the percentile distribution charts (Chart 10) you can see that Strategy #173 is better than #89 all the way down through the 25th percentile. The IRR at the 25th is 8.65% for Strategy #173, which is slightly better than 8.54% for Strategy #89.
It depends on management's assessment of these results and its risk tolerance. If they want to win 75% of the time, then they may be willing to go to Strategy #173. That is the type of analysis we do with this graph.

CHART 10
DISTRIBUTION OF RETURNS

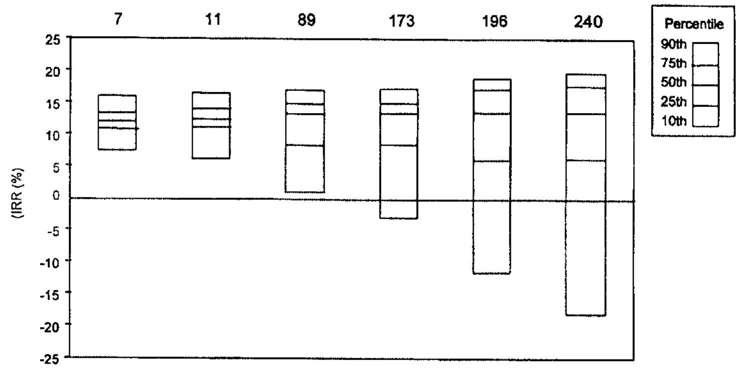

Distribution of Results (Percentages)

Percentile	7	11	89	173	196	240
90th	16.19	16.68	17.17	17.42	19.03	19.89
75th	13.78	14.08	15.12	15.18	17.35	17.83
50th	12.27	12.69	13.48	13.64	13.75	13.77
25th	10.87	11.07	8.54	8.65	6.28	6.46
10th	7.49	6.31	1.20	(2.85)	(11.37)	(17.83)

A lot of time should be spent in the fourth step, which is sensitivity testing. All the assumptions used up to this point have been the base-case, best-guess assumptions. Now I need to understand how the key factors drive the results. I want to know how changes in any of those key factors would affect my decision. How would alternate assumptions affect the key frontier? Key factors could include lapse rates, RBC targets, competitor-credited rates, crediting strategy, and capital market variables such as the shape of the initial yield curve or the market yield spreads. I use the results of sensitivity testing to refine my optimization.

I need to mention that, when I do sensitivity testing, I may continue to test strategies that are not on the efficient frontier under the base case (Chart 9). I may keep strategies that are reasonably "close to" the frontier because changing the assumptions may cause some of them to move up to the frontier. Obviously, I would want to know that information.

Chart 11 shows what might happen if you change your base assumptions. One line, EF_1, represents the efficient frontier from our initial base-case assumptions. It is relatively flat. If I change assumptions and discover that the shape of the efficient frontier changes, then I may decide to move out further on the efficient frontier because the extra reward may offset the additional amount of risk. Chart 11 shows that we may prefer strategy #1 under the original assumption set, but we may prefer strategy #2 under a different set of assumptions. In the near term, we will focus our target mix on strategy #1. But, as factors change and as experience develops, we will keep in mind the characteristics of strategy #2 so that we can be making adjustments to our portfolio. This approach also gives our portfolio managers some room to take advantage of trading opportunities.

CHART 11
CHANGING BASE ASSUMPTION

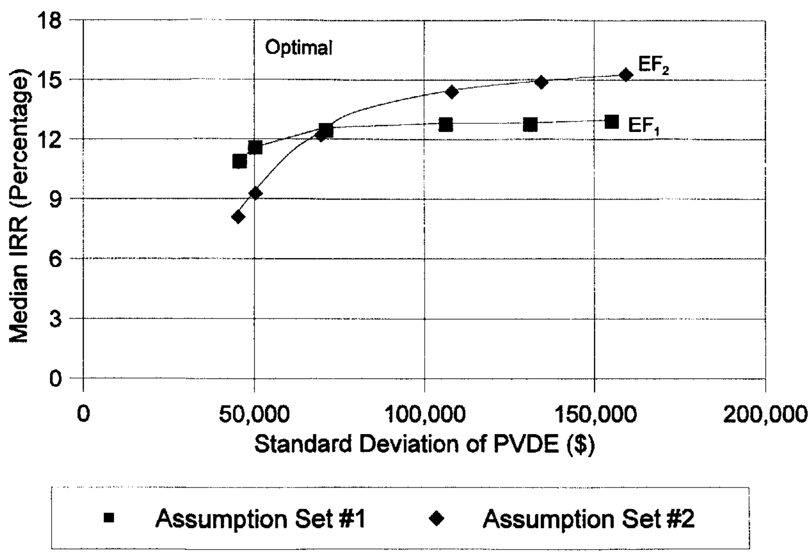

Asset Allocation:	#1	#2	Acceptable Range
3 year	40%	20%	20-40%
5 year	40	30	30-40
7 year	20	40	20-40
10 year	0	10	0-10

An example of an assumption change that might cause the efficient frontier to steepen is a change is the crediting strategy. For the case study, I changes the crediting strategy to hold the first year credited rate constant at 7.5% for all strategies. This means the year spread differs between the different investment strategies. You can see that with shorter strategies, I am actually reducing the target spread because I am declaring a higher credited rate, but in the longer strategies I am picking up interest spread. You can see in Chart 12 that now the longer strategies may pay off. I may now feel comfortable going with longer assets, as long as I keep the extra spread available from those longer strategies.

CHART 12
CREDITING STRATEGY CHANGES

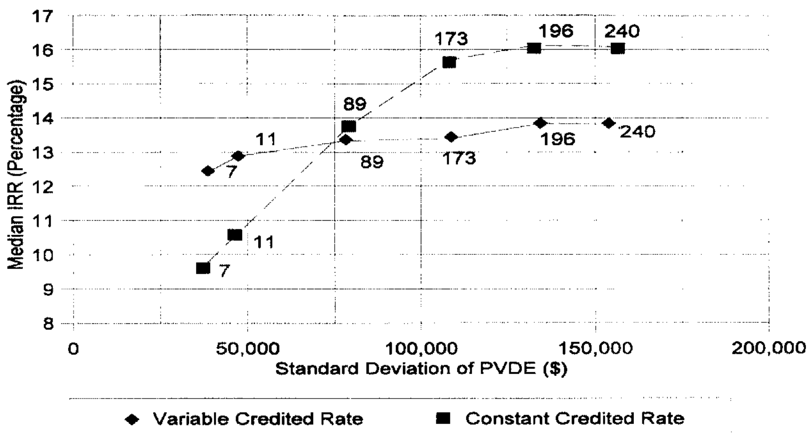

Variable Credited Rate	7	11	89	173	196	240
FY Target Spread	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
FY Credited Rate	7.38	7.42	7.65	7.71	7.77	7.84
Constant Credited Rate	**7**	**11**	**89**	**173**	**196**	**240**
FY Target Spread	-0.12%	-0.08%	0.15%	0.21%	0.27%	0.34%
FY Credited Rate	7.50	7.50	7.50	7.50	7.50	7.50

A similar result is obtained if you are testing riskier asset classes. Many times the marketing people say that they need a higher credited rate to be competitive. Therefore, they ask us to go out and take more asset risk so that they can get a higher portfolio yield and can afford to pay a higher credited rate. To combat that mindset, I show them an efficient frontier graph to illustrate that, if we take slightly more risk, the nominal portfolio yield will increase. However, you may not be able to credit a higher rate. You may need to keep all the extra spread just to match the less risky strategy that we were comfortable with, such as Strategy #89, in the first place.

Another sensitivity test that I did for this case study was to increase the sensitivity of lapses by doubling the lapse function. You can see in Table 6 that you do take a hit at all the different percentiles. For Strategy #89, the median return has decreased by about 1% from around 13.5% to 12.5%, but the 10th percentile has declined by close to 5%. This fact is important if your management is interested in looking at downside risk to make its decisions.

TABLE 6
INCREASING THE SENSITIVITY OF LAPSES

	Strategy 7		Strategy 89		Strategy 240	
Percentile	Original Lapse Assumptions	Double Base Lapse Assumptions	Original Lapse Assumptions	Double Base Lapse Assumptions	Original Lapse Assumptions	Double Base Lapse Assumptions
90th	16.90	15.99	17.17	16.97	19.89	19.69
50th	12.27	11.37	13.48	12.50	13.77	12.81
10th	7.49	4.49	1.20	−3.80	−17.83	−24.87

If you were to graph this in a mean-variance framework, you would see that this change would have simply shifted the efficient frontier downward in parallel. I probably would have still picked strategy #89 because the curve's slope will flatten at the same place as before, but management needs to understand that 10% of the time it will still experience very negative returns. If that is not acceptable, then you may need to move down from strategy #89 to one of the less risky asset portfolios.

The final step is portfolio selection. Based on the base-case results from step #3 and the sensitivity tests from step #4, we will identify where we want to be on the efficient frontier. We typically will identify a range of portfolios on the efficient frontier with which we are comfortable. We can build in any management biases at this point. To the extent that your company has an interest rate forecast, you can use information to decide which end of the range to move toward. You may choose to select a strategy that includes any asset classes in which you are currently overweight. Although those asset classes may make sense in this analysis, you may not be able to move into them at that time. You also need to take into account any existing portfolio mismatch. To the extent that your overall portfolio is too long for your liability mix, you may choose to put new money to work according to shorter strategies.

Based on the sensitivity analyses and current market conditions, we recommend an investment portfolio strategy. We choose a current target mix, which would be our short-term strategy. We also select a portfolio range based on the sensitivity tests. Note that, if you do the sensitivity testing now, you will save a great deal of time later on, and you will better understand the relationships between your assets and liabilities. As conditions change and experience develops, you can change your target mix. The key assumptions underlying a product strategy might be the yield curve shape or market yield spreads. Market spreads may be tight currently. But if spreads widen in the future, you may decide to move out of one asset class and into another. This information is developed at the sensitivity testing stage and then used to develop the portfolio range and a target mix.

In conclusion, I have described Falcon Asset Management's approach to a totally integrated asset/liability management process. Remember that the liabilities drive the assets. Because we are making decisions under uncertainty, we feel that 100–500 stochastic simulations are needed to generate a wide range of possible returns to allow management to understand and manage the risk-versus-reward trade-offs. Sensitivity testing is a key part to the ALM process. You need to work closely with the product managers and product line actuaries to set the assumptions. By joining the investment and product line actuaries in the asset/liability management process, you create better solutions for both sides of the balance sheet. Asset decisions include the choice of asset classes and the management of asset durations. Liability decisions include better product design. Also, you need to determine the target spreads appropriate for the level of risk you take. Each insurer's portfolio should be tailored to meet its specific product line characteristics. Each company has its own distribution system and target market. Therefore, there is no one set of assumptions that is right or wrong. This is why sensitivity testing is so critical. Finally, ALM is a continuous improvement process. You need to test many assumptions, analyze the results, and try to understand why things are happening. Then you need to build on what you have learned so that the next time you do the project you can incorporate that information and do a better job.

MR. GEORGE: Let me just make one or two follow-up comments, and then we'll open it up for questions. This type of analysis has limitations as our models are not perfect. Brian pointed out that they're based on assumptions of which we're unsure. Certainly, policyholder behavior and agent behavior can change the results that come out of our analysis, as Brian also pointed out.

Our models can be limited due to the fact that they lock us into certain strategies and behaviors upfront that occur under all interest rate scenarios—they don't really mimic how our companies would behave as conditions change. Because of that, this type of analysis doesn't really give all the answers. It provides very good insight, and I think it's necessary, but you wouldn't want to use it to give you all the answers. You can't do it once and then be done. It's the kind of thing that you need to do periodically. You want to do it often; some people do it monthly. Depending on the risk involved in your portfolios, you might do it monthly, quarterly, maybe semiannually at a minimum, but it depends on the nature of your products as to how often you would want to do this. But you don't want to just do it once and leave it; it needs to be done periodically. With that, let's take questions.

FROM THE FLOOR: If I understand your process, to develop your efficient frontier you developed more than 250 unique investment strategies and then ran each one of those through 100 scenarios. Is that the process?

MR. TRUST: Yes, for new business models. For this example, I ran 286 different investment strategies. Realistically, though, when I'm running in-force models there's not enough time to run 286, so I will cut back the number tested. I've used Doug's approach to try to get an idea of what kind of asset duration I might need, and I will try to limit my search area to a much narrower range. When I am running new business, though, I often have no idea where to start. Running a large number of strategies is a good place to start. That is what I was hoping to communicate through the efficient frontier graph with the "blob" of dots (Chart 9).

FROM THE FLOOR: I have a second question. I assume when developing your IRR, using discounted distributable earnings, you're going to run into cases in which your IRR solution isn't unique. How did you treat that?

MR. TRUST: That is why I use the median IRR instead of the mean as the reward measure. By looking at the median, I am throwing out the outlier scenarios. I am looking for the central tendency based on a median instead of a mean.

MR. JOSEPH L. DUNN: The thought occurred to me when you were discussing liabilities that although you are trying to model the interest rate sensitivity of the policyholder option, one aspect is being left off, and that is that the policyholders may exercise that option based on their judgment of the credit quality of the company. In other words, you are subject to a run on the bank, usually at the time when you can least afford it. That seems to be a major aspect of risk analysis. Have you ever given any consideration to including that type of behavior in your models?

MR. GEORGE: Yes, I think that's a very good point. I'm not sure if it's a matter of dynamically including the behavior or maybe just creating a few scenarios to look at when this type of behavior occurs. Let's look at a worst-case scenario. Once we get to the point where we're narrowing down our decision, let's look at a worst-case scenario. If we do the rebalancing, what happens during a run on the bank? When it comes down to it, that's why we're doing this analysis. We're very much afraid of extreme conditions. For most interest rate changes you're not going to have to sell many assets. You're not going to have to change your risk profile to handle most conditions. It comes into play under really adverse conditions, and that's where this kind of analysis will benefit you. Many companies haven't done this kind of analysis during the last 10–12 years, and maybe they got lucky because interest rates have been dropping so much, and they were probably too long based on our analysis here, and they made money because interest rates kept going down. That's a nice, safe environment, and if you could continue to count on that, then you wouldn't need this analysis. We're very afraid of interest rates going back up and hitting those levels again, just like they did in the early 1980s. That's what's going to produce this run-on-the-bank scenario. I think that's a very good point—do some real stress-testing once you get the alternatives laid out and see how good it works. Let's put our solution through some real stress-testing scenarios and see if this is really helping me.

MR. TRUST: My ultimate actuarial model will some day include the capability to model all the assumptions that go into the model as dynamic assumptions, including some kind of stochastic process instead of static assumptions. Perhaps that approach could build in a run-on-the-bank tied to the lapse assumption.

MR. TERRENCE M. OWENS: Brian, was that a single policy that was modeled or a portfolio of SPDAs?

MR. TRUST: It is a block of new SPDA policies representing $10 million in premium issued at a single point in time.

MR. OWENS: You indicated the portfolio strategies that you were looking at were investments of cash flow. For example, if you had a 20% allocation to three-year assets, would only new cash flow be allocated 20% into three-year, or would that 20% of the investments be kept in three-year assets at all times?

MR. TRUST: Positive cash flows would be invested 20% in three-year assets. No rebalancing is done. As positive cash flow occurs, the money is put to work according to the initial mix. Most of the cash flow occurs upfront anyway on a SPDA product. A rebalancing-type effort will not really impact the results as much as you might think.

MR. OWENS: Douglas, you indicated that the 10-year negative key rate duration was due to the crediting strategy. Was one of the key rate durations, in fact, negative?

MR. GEORGE: It's slightly tricky, but when you get out on the yield curve, two factors occur. One is you're past the surrender charge period, you're going to see very interest-sensitive cash flows coming out of your model, and that'll tend to bring down duration, or key rate duration in this case. That tends to bring it down because we're past the surrender charge period. The second factor that tends to bring it down is because our interest-credited rates are based on points further in the yield curve. In this case we're not using the ten-year rate. It's, say, a five-year rate, which is more typical for an SPDA. If you're out five years and you're crediting based on the five-year rate at the time, then that ten-year rate increase will increase the cost of your SPDA.

MR. OWENS: I guess that brings up my question to focus then. If you use the five-year crediting rate, would you expect to see little change to the four-year key rate duration?

MR. GEORGE: That's a good point. You need to consider, when doing this key rate duration analysis, how to set up your model. We set up our model to base it right on the five-year rate. That's not the way we're going to credit interest. If the five-year rate were to shoot up 100 basis points, and the rest of the rates were to stay the same (I realize this is unrealistic), would we truly credit 100 basis points higher?

That's what happens when you do the key rate duration analysis—you need to consider the way you set up your model as to the effects that has on how the key rate durations come out. That's one of the reasons that I advocate not chopping up the yield curve into too many pieces. I've seen people do key rate duration analysis and chop it up into 11 key rate shifts. I think it's overkill. If my 20-year key rate duration is off one way, and the 25-year key rate duration is off the other way, I don't think I'm too concerned about that because those two points on the yield curve aren't moving independently.

MR. RICHARD JUNKER: I was interested in the mention of assumption for inflation tied to Treasury rates in the short term. That would have been a good idea, but then during the last few years we've seen that despite major interest rate movements, inflation has been very stable. Is it truly a very good time to use short-term rates?

MR. TRUST: I haven't performed any studies on that subject. Inflation is very important in the analysis for some products, such as universal and traditional life. It is not as important an issue when it comes to analysis for an SPDA product because typically the expenses have relatively little impact. Only use inflation in modeling future maintenance expenses. We must examine all the assumptions and determine the appropriate point on the yield curve to tie the inflation rate. It may turn out to be a function of several points on the yield curve.