2005 Valuation Actuary Symposium*

Orlando, Fla.
September 22–23, 2005

Session 7TS
Hedging and Other Mitigation Techniques

Instructors: Ellen Cooper
Andrew Rallis
Ken Mungan

Summary: This session describes current risk mitigation techniques and their impact upon valuation systems. Topics include hedging mechanisms, regulatory relief, reinsurance and securitization, product control cycle, realistic assumption changes, maximizing returns and minimizing litigation. At the conclusion, participants understand common risk mitigation techniques and financial reporting issues related to them.

MS. ELLEN COOPER: I'm going to start by introducing our panel. Our first speaker is Andy Rallis. He is a senior actuary from Metropolitan Life Insurance Co. His responsibilities include asset/liability modeling (ALM) and variable annuity (VA) hedging. He graduated from MIT. I'm a partner from Ernst & Young, LLP, and I lead up our development initiatives regarding risk and capital management, specifically regarding VA hedging. Our third speaker is Ken Mungan. He is a principal at Milliman, Inc. He's the leader of the financial-risk-management practice. He also graduated from MIT.

MR. ANDREW RALLIS: I'm going to give you the Metropolitan Life Insurance Co. perspective on VA hedging. Metropolitan Life put its hedging program in place in July 2004. It covers living benefits. The company had been issuing living benefits for about two years prior to putting the hedging program in place. I'm going to go through the main lessons that we've learned. This is our top 10 list of things that matter from the perspective of actually having done the hedging.

1. Accounting

Note: All charts referred to in the manuscript can be downloaded at:
http://handouts.soa.org/conted/cearchive/valact05/2005valact_handouts.htm.

2. Software
3. Hardware
4. Communication
5. Accountability
6. Delta
7. Interest Rates
8. Volatility
9. Behavior
10. Risk Profile

I'm going to go through these items in my presentation, but it is interesting thing to note that, from my perspective, the operational risks are at least as important and probably more important than the actual financials. They drive a lot.

With accounting, we're trying to balance three different perspectives of risk. The economic perspective of risk: What would the capital markets say about these things? The GAAP perspective: What does U.S. GAAP say about these? And the statutory perspective: What do the regulators say about these things?

From an economic point of view, we're concerned about the economic gain or loss to the enterprise. How much is the enterprise really worth with these guarantees in place, and how much will that fluctuate over periods of time? From the GAAP perspective, we're more interested in income volatility. How did GAAP earnings emerge under these conditions? With statutory, the concern is usually solvency. Is the company going to be around in the future? One lesson is that small product variations can drive a big change in strategy because of how these things get classified under GAAP and modeled under statutory regulations.

There are two popular forms of guaranteed living benefits: the guaranteed minimum income benefits (GMIBs) and the guaranteed minimum withdrawal benefits (GMWBs). Guaranteed minimum account balances or accumulation benefits (GMABs) are emerging. Like all good products, the nature of the benefit is very easy to describe. From the policyholder's point of view, it's a very easy thing to envision. GMIBs guarantee income for life. GMWBs guarantee return of principal. There will be additional options on those, but that's the main theme. They are very easy benefits to understand from that point of view.

GMIBs are classed under GAAP as insurance contracts and GMWBs as embedded derivatives. For GMIBs, that creates an accounting mismatch, because any instruments that you use to hedge are going to be marked to market as derivatives. The hedging instruments are not insurance contracts, so there's a mismatch for hedging that is created. With GMWBs, the hedging instruments are derivatives, the embedded derivative is a derivative, and there's an accounting match. And that makes it substantially easier to design a hedging program for GMWBs.

For guidance, the relevant accounting literature under GAAP includes Standard of Practice (SOP) 03-1 for GMIBs and Financial Accounting Standard (FAS) 133 for GMWBs. You can look to these particular paragraphs and DIGs for guidance on the key issue of net settlement. Can the policyholder settle the contract in cash? If they can, it's considered embedded derivative. If they can't, it's considered insurance contract. Those are two very different accounting paradigms. If you are classed as an insurance contract, your assumptions need to be consistent with the assumptions that you use in your deferred acquisition cost (DAC) model. That's important. But if you're classed as a GMWB, your assumptions are consistent with capital markets. For both products, you're going to have similar types of behavior assumptions; the key difference is in the interest rate assumptions or the modeling assumptions for economic scenarios. Your DAC assumptions are more likely to be based on a realistic interest rate; your capital markets assumptions are going to be based on risk-neutral interest rates.

Software is the next most important thing. Once you've figured out your hedging strategy, you've got to be able to do a tremendous volume of calculations in production. We are doing these calculations daily, overnight. Speed of execution is important, because we're creating millions and millions of scenarios. We're doing policy-level scenario analysis. Right now, those run overnight on a bank of 64 servers. And it still takes seven hours to run these calculations. Connectivity is important. It goes along with production environment. How do you get information in and out of the model? We're connected to Bloomberg terminals so that we can update capital market assumptions and stay current. Support is important. What were we doing two years before we were doing hedging? We had a lot of internal models. The internal models were very good for pricing purposes and trying to understand our risk. But the real question is, are they suitable for making day-to-day decisions? That's why we needed to create an environment in which we were using third-party software to do that.

Consistency of hedge and reserve calculations—on a capital-markets basis, we have to do certain calculations, like our GMIB. We have to do SOP 03-1 calculations. For the GMWB, we have to do FAS 133 calculations. We wanted all of those calculations to come out of the same system. We didn't want to have one system for hedging and a different system for reserving. The last consideration is stochastic-on-stochastic. This goes along with a pricing analysis or a risk type of analysis. We didn't need to do this in production. But with the current risk-based capital (RBC) requirements, we're pushed in the direction of stochastic-on-stochastic models as part of the production environment. To get credit for the hedging program, you really have to model the impact of the hedging program over each scenario, which requires stochastic-on-stochastic.

Hardware matters in the production environment. We have a test environment in which we test out software. Then we have a production environment that is isolated from everything else. There's also a fail-prevention environment. If our production servers go down, we can switch to a different set of servers that are reserved in

case of a disaster. We take the production environment seriously. A lot of companies have distributed processing. This application lends itself to distributed processing. You can parse out the calculation and do it on different computers that are operating simultaneously and then assemble the results. We've got dedicated servers. Other insurance companies have distributed the application over their desktops. We don't do that. We don't think that it's really a good idea when it comes to security. We have an isolated environment in which just this application is running on those computers. We've got 64 servers, and we're going to increase that to more than 300 servers probably this year. Speed of execution is very important. You have to do these calculations every day, so you need enough computing power to be able to redo the cycle within 24 hours at the most. Right now, we're taking seven hours.

The classic Greeks are measured as part of a hedging program—option value, delta, rho and gamma. From a risk point of view, these are virtually indistinguishable risks for both GMIBs and GMWBs. They have the same profiles, but the accounting paradigm drives you in two completely different directions.

Why does communication matter? There are a lot of constituents here that want to understand the hedging program. We spend a lot of time describing what we're doing to all of these constituents. We have boards of directors. I'm using the plural because, in our company, we have multiple legal entities that are involved. Some of them have their own external directors. They obviously have an interest, too. We spend lot of time in front of boards of directors explaining the program and what we're doing. We have our own senior-management team, the executive group, our internal corporate risk-management area, treasury and valuation. Obviously, there are complex issues involved for the cash-management issues and capital-management issues, and valuation, obviously, is getting more involved with the C-3 Phase II work.

How do you audit a program like this? This is a major undertaking. We have external constituents, the regulators and the rating agencies. We have a once-a-year meeting with the rating agencies, at which we explain the hedging program. It's hard to do in the limited amount of time that we have. We've had multiple meetings with regulators, trying to get them comfortable, as well.

Part of communication is reporting. We have developed a reporting template that we calculate and circulate once a month to a number of people in management. We circulate it to about 10 or 15 people so that they can keep abreast of what's going on. That reporting template includes an attribution piece. If we have a change in value, how much did the value change from the previous period (the previous month-end to this month-end)? We also include the Greeks.

When you're trying to communicate such complex topics, it’s helpful to use a graphic. A simple graphic like the one in Rallis Slide 11 is easier for people in senior management to understand than if I had simply shown six numbers on a piece of

paper. It communicates that we have a certain amount of economic value from the beginning of the period. It has changed by the end of the period. These are the components of the change. It gives you a sense of the size of each of those changes. Don't skimp on the time spent to create this kind of graphic. This took many hours of computer time to generate, but in order to tell the story you need really compelling graphics. Otherwise, people's eyes glaze over.

Accountability was one of the first things that we had to address with internal auditing. Are people's roles clearly defined? Running the hedging program effectively requires the attention and contribution of many different departments of the company. The roles have to be clear as to who is responsible for what. We, in the ALM unit, act as the program sponsors. We recommend the hedging strategy. We oversee the day-to-day calculations. We are responsible for the implementation of it, making sure that the hardware and software are in place. But there are many other departments that have a stake. The product-development area is closely involved. We have to have very tight communication with them about their product plans or their plans for when they change funds in the VA offerings. We have some reinsurance in place with outside reinsurers, as well as captive reinsurers, and we make sure that those treaties are all in good shape. There is a lot of interaction with accounting on SOX issues, getting the numbers into the ledger. For the technical accounting unit, we classify these contracts—whether they were SOP or FAS 133. We also communicate with valuation, the tax department, the finance department, risk management, auditing, etc. We are in daily contact with the investment department. We've set the guidelines. They execute the trades to keep us hedged within the guidelines. Then we monitor that activity on a day-to-day basis. And there is probably a list of 30 people from IT that are involved in this.

Delta is a first-order change in equity. It matters for the obvious reasons. It affects the economic value. One consideration is that SOP liabilities react to it less sensitively than economic or capital-market types of assumptions. A big consideration in our program design is balancing the hedging of the economic value against hedging GAAP income volatility. We've made some trade-offs that we're comfortable with, but it involves less than full hedging of the economic value for our GMIBs and essentially full hedging of economic value for our GMWBs. But even in the case of GMIBs, when we're not fully hedging the economic value, we're putting catastrophic protection in place, as well.

Interest rates matter a lot, probably surprisingly. The rho calculation is sensitive to different things. Interest rates enter this calculation in different ways. They enter through the discounting of claims and fees. They enter through the modeling of the bond-fund performance. And in the case of GMIB, they actually enter into the level of claims because the GMIB benefits are driven by how close you are. GMIB, basically, gives the right to buy an annuity stream at a guaranteed predetermined level. And there is an interest rate that that annuity stream is based on. So GMIB locks in that guarantee. It also drives the level of claims.

Rallis Slide 18 shows that cross-Greeks are important, particularly in the impact that interest rates actually have on the hedging program that you have in place if you're hedging equity. That's surprising, too. It has a large impact on the amount of equity risk. Think of it this way. In the case of GMIB, if we were far out of the money with respect to interest rates, the guarantee in the annuity phase may call for a certain mortality table with 2.5 or 3 percent interest. If interest rates are 7 or 8 percent—well above that guarantee—your equity risk also disappears, because you can't have one risk without the other. This slide shows the magnitude of these cross-Greeks. This is the impact that yield-curve changes have on delta itself—not on the value of the option that we've written, but on the delta of the option that we've written. If you're following pure delta hedging, the cross-Greek would be flat. There would be no cross-Greek if you were doing only delta hedging. But the cross-Greek in the liability is rather large. It's even more pronounced than the cross-Greek of a put option. If you think that you can be hedged fully even with a put option, you can't. Because as soon as interest rates move, you're going to have to adjust your position.

Volatility is measured through vega, the sensitivity of the economic value to equity volatility movements. The SOP is not as sensitive to these movements as the economic value is. Marked to market, you have to look at implied volatility. With the SOP calculation, there's a little bit more latitude. And you've got to make a trade-off between looking at implied volatility and historic volatility. But at the end of the day, it's likely that your SOP calculation is going to be less sensitive to vega than your economic-value calculation. Again, you've got to make that trade-off.

Vega, as we said, matters a lot, but it's hard to know where to hedge vega. There's vega embedded in any derivative that you are going to buy. We write the GMIB. It's got a 10-year waiting period; it can be executed any time thereafter. The issue arises: What kind of option would you hedge it with? How much vega do you really need to protect? Where along the volatility surface should you be? We did a calculation of key-rate vega and found that no matter what the "in-the-moneyness" was, we'd be best protected if we concentrated vega at around five to seven years. A lot of that has to do with the shock-lapse rates at the end of the surrender-charge period; most of the vega is concentrated before that.

Why does behavior matter? We have very little experience in the U.S., and especially in the GMIBs, in calibrating the parameters of the behavior models. Essentially, behavior risk is unhedgeable; there are no capital-markets hedges for behavior itself. We know that we're not going to be able to hedge it. We've developed a model for the residual-behavior risk. We've come up with an alternate set of assumptions for the behavior model, the lapse rates and the annuitization rates. What is the probability that you could switch into a very bad regime of behavior, like if everybody acted in concert because they were reading the same *Wall Street Journal* article? How bad could it get? We've developed a probability that that could happen. Based on that, we've come up with additional internal-

capital factors of 1 to 2 percent, basically unhedged, of economic capital that we would need to support these risks.

Risk profile matters because of the diversification offsets. We know that these riders add a lot of equity risk to us. We've seen that they could add a lot of interest-rate risk to us as well. Do you like that risk or not? Insurance companies get paid for taking risk. You have to try to understand how that risk fits in with the rest of the company. We look at our rho in the context of the rest of the company's interest-rate risk. I'm responsible for all of ALM, including the general-account ALM. The rho becomes part of our duration match or mismatch position. With the equity component, you've got to consider (at least in the C-3 framework) that it can be impacted by the other elements of your RBC calculation. If you have a lot of credit risk in the company, you could be diversified with respect to the equity risk. This is maybe something to consider.

Your RBC ratio could depend on the other risks in the company. If you have a choice between different legal entities, this may be something that you want to look into. If you put it in a company that's just issuing variable annuities, you're not going to get much diversification offset. If you're putting it in a company that issues a lot of other insurance contracts, you'll get more diversification offset. Our internal economic-capital formula is dominated more by credit risk and, to a lesser extent, by liability optionality and interest-rate risk. We found that we have a higher tolerance for equity-risk exposure than we might have otherwise. It gives us some comfort, especially in the case of the GMIBs, that retaining some of that equity risk is actually a good thing for us and that we're less inclined to try to get full protection on it.

MS. COOPER: I'm going to give an overview of some of the issues that surround hedging and focus on hedge-strategy simulation. Mr. Rallis talked about overnight runs, production and implementing a hedge strategy. We are often asked questions about what happens over time. I know that I have tools now that I can use to run in production and get liability Greeks and a hedge strategy for today, but how do I know over time what the cost of that strategy is? How do I know over time what my economic exposure is, what my income volatility is, what my capital volatility is? These are some of the questions that our clients are asking us. We have tools that can help us to answer these questions.

We're going to talk about a GMWB today, because we can focus on an embedded derivative that is FAS 133. I can show you a hedge strategy that uses the same type of accounting, and we don't have to deal with the issues of GAAP income volatility due to accounting mismatch. Having said that, Cooper Slide 4 shows an example of a new business GMWB with $1 million of initial premium and a benefit cost of 50 basis points. We're going to assume that we can take out a maximum of 7 percent per year. There's a 135-basis-point charge on the base mortality and expense and a split of 60/40 between equity, and bonds and money market.

One of the first things that we have to do when we look at the liability structure is, once we have our product features, we need to come up with our appropriate assumptions. Throughout the industry, this is a big question. I'm going to show you, on a later illustration, how sensitive the hedge cost is to some of these assumptions. I recommend that you understand how varying these assumptions are, because at this point we don't have experience. These are true assumptions, how sensitive they are. In Cooper Slide 6, the lines representing pure income buyers. These are the contract holders that we believe are only going to elect when their benefit is in the money, when capital markets go down. If the markets are moving up steadily over time, not very many people are going to use this, because we believe that people are buying this as a protection when the markets go down.

In opposition to that, is the line that represents a protection buyer. They're looking for a steady flow of income and are using the GMWB to protect their income, 7 percent per year over a number of years. What happens there? In a "steady up" environment, you're going to see that there's going to be significant usage still. If you look at "steady down," you're going to wind up in the same place. We need to think about that. Some of the other items that we need to think about include dynamic lapse and mortality. In the case of a living benefit, people dying, unfortunately, helps our exposure to the living benefit. Finally, our capital-market assumptions are critical to our analysis.

When we're dealing with a hedge-strategy simulation, not only do we need to think about our risk-neutral parameters at time zero, but we need to think about how our scenarios move through time. How do we illustrate that so that we can have an appropriate measure? There are many different ways in which to measure that. In the example in Cooper Slide 8, we hedge strategy simulation. For simplicity, we've assumed an at-the-money, implied-volatility structure. We move up that curve, moving up the forward rates, as we move through time, similarly to how you would do this for interest rates.

Before we go into hedge-strategy simulation, let's look at what happens at time zero under varying assumptions. Our assumptions can allow our time-zero liability cost or our hedge cost at time zero to vary by up to four times. I can construct a scenario for which, if I believe that I'm going to allow my policyholders to drift their asset allocations, and I don't believe in skew (which I don't recommend, but this is just for illustration purposes), I can show you a liability cost of eight basis points.

Now, let's move to Run 2. Run 2 assumes that there is a rebalancing in place. Think about what that means. If I have 60 percent equity, 40 percent bond at time zero, in the next period, if I'm in a down cycle, I'm going to rebalance so that I have 60 percent equity. I'm buying into a down market. That's going to make it more expensive to hedge, but it's also going to make it easier to hedge. So there are trade-offs. It is easier to hedge, because you know what the asset allocation is moving forward.

In Run 3, we've added skew, still at time zero. What does skew do? With skew, when volatility is going down, I'm going to expect my volatility to increase. That's typically what we observe in the market, although not all companies have, in fact, incorporated this into their strategies. You can see that it, again, increases the cost at time zero.

Finally, if we assume that everybody uses this option, as in Run 4—if we assume that this is really a protection-buy on income—our liability cost goes up to 32 basis points. I recommend that we understand what the different bells and whistles do to our cost and our results and that we're able to illustrate this to senior management, to the board, to all of the constituents, so that they understand that any result that we show is not the "answer at the back of the book." It's really an interval. There's a range of what can happen.

Now we're going to examine the world of hedge-strategy simulation. The first time that we tried to do a hedge-strategy simulation, the client wanted to see how this looked through time. That was probably two years ago. At the time, the tools weren't accessible to us. To run a single cell through time, it took 42 hours. With today's technology, we can now run this in about two minutes on a single machine. We went from being able to go home for the weekend while doing one of these, to not even being able to run to the bathroom before you have your results. This really enables us to get to the answers to questions and do analysis that we just didn't have at our fingertips before.

Cooper Slide 9 shows path-wise liability value under alternate paths. I'm going to talk about Scenario 1 on the slide. This is an experience scenario in which we're going to assume that the market goes down 10 percent per year every year. That's a bad scenario. What does that mean at the end of 10 years? It means that there's nothing left. You can see what happens to my liability value. This calculation is the present value (PV) of claims minus PV of premiums. You can see what happens over time. This scenario goes way in the money and ends up with this huge liability value.

Let's look at our other Greeks. Cooper Slide 10 shows delta, the exposure to equity. You can see that it peaks at five years. After that, this is definitely pointing to being significantly in the money. By the time we get to the 10-year period, we're not helping ourselves by being exposed to delta. We're already significantly in the money, and there's almost nothing left that needs to be hedged. There's no risk left. It's a certainty that we have a claim.

Cooper Slide 11 shows vega and rho. Scenario 1 shows Standard & Poor's (S&P) assumption, linear down 10 percent per year. We rebalance across time. At time zero, I'm doing thousands of scenarios to calculate Greeks. I'm calculating my liability fair value, my delta, my gamma, my vega, my rho. Then I move through time assuming 10 percent per year. This is just a very simple illustration. We can do this stochastically. We need to do this for C-3 Phase II, for those of you that

have hedge-strategy in place. At the end of that first year, the S&P is down 10 percent. From where we are in our experience, we regenerate, on the fly, thousands of stochastic scenarios. We recalculate liability value, delta, gamma, rho and vega. We look at the hedge strategy that was in place at time zero and make the appropriate rebalancing to reflect our match.

Vega is an indicator of certainty or uncertainty of payoff. It peaks at three years and falls down at five, because by five years there's almost a complete certainty that there's a payoff. And rho, because it's a GMWB, follows a similar pattern to delta.

For simplicity, again, Cooper Slide 12 assumes a hedge strategy on top of this. The hedge strategy that we're going to assume is a very simple 10-year, at-the-money, European option. We're going to rebalance over time with that option to vega. Then we're going to fill in with delta and rho using equity futures, using interest-rate futures. As we move through time and we look at our economic numbers, what do we see? Over that period of time, with S&P linear down 10 percent per year, we've got this huge amount of claims. The PV of claims back to time zero shows $127.5 million worth of claims.

What's interesting is that we've hedged our premium. If you looked at the PV of the premium at time zero, you would see that there was an expectation. At time zero, when everything was at the money, you would have approximately $40 million of premium. Our hedge worked, because we were able to save that. We hedged using futures or premium, and we have close to $40 million. It's showing us that the hedge works.

We have an initial outlay at time zero. The option payoff essentially is reflecting over time the continuing rebalancing, the continuing purchase of European puts, and then, finally, payoffs. We decided to use 10 years. But 10 years is probably too late, isn't it, because what happens at the end of 10 years? There's no account value left. Just for illustrative purposes, we show that you still end up with a total accumulated cash flow net of the initial outlay of $125.8 million. And with claims futures, because there is delta in the option, we actually end up slightly overhedging in the future strategy. We wind up with a negative in terms of claims futures PV.

We end up with a net-hedge cost of 47 basis points in this awful, dismal strategy. The initial premium was 50. We wind up, even in this situation, being able to demonstrate that over time in this awful situation with a hedge strategy that is not necessarily our most optimal, we still end up basically breaking even. This is something that could be a useful tool to show to senior management that this really is risk-free for us.

I don't particularly think that the strategy that we showed was a good one. What are some of the other things that we could do? We could look at alternatives for

some out-of-the-money options. We also want to look at some strategies that have a basket with tenures that are lower than 10 years, so that when we start hitting our first claims in this dismal strategy, we actually have some cash flow coming from the options. Why is this not a great strategy? Well, 10-year options can be somewhat illiquid and expensive. We don't want to rely on them. We already talked about having payouts of claims in these dismal strategies. It's directional. We're buying in a down market.

There are many reasons why this is a con. I recommend that whenever you're looking at a hedge strategy, not only do you look at it over time, but understand your alternatives and the pros and cons of different methods. Be able to demonstrate them to senior management, to your board, to your internal constituents. In Cooper Slide 14, the solid line represents the amount of notional needed to maintain the hedge at 100 percent at the money. The dotted line shows the amount required if I went to 70 percent. You can see that, although I have some exposure, I end up needing to hedge significantly less. I'm going to have significantly less outlay, significantly less notional value.

Why is hedge-strategy simulation important? We need to look at three pieces together. I focused on the economic piece. What does the hedge cost? There are two other really important pieces that need to be balanced with this. One involves GAAP income volatility and understanding how your income is going to change from period to period with the hedge strategy that you've chosen (and the alternatives) and your capital. We're faced with C-3 Phase II. We want to put our dynamic-rebalancing strategies into our capital calculations. We want to be able to do a capital-even simulation and understand how our capital is going to change under C-3 Phase II from one period to the next. We need to understand all three of these—what we call the triangle—in conjunction with each other, in order to appropriately analyze our hedge strategies.

MR. KEN MUNGAN: As we start going through the material on hedging, one point I want to make is that for effectiveness in an enterprise-scale hedging program of this nature, it requires a very collaborative effort. I've seen that firsthand at many client companies and within Milliman as we embarked on this adventure to help companies with hedging. I recognized very quickly that my own skills as an actuary could solve only part of the equation. We needed to assemble a broad group of people that included actuaries, traditional Wall Street professionals, capital-markets people familiar with trading, and an entire software-engineering team that brings technology expertise to the table that goes beyond the traditional technology used in the life-insurance industry.

We need to understand how hedge programs perform over time. As we started working with major investment banks on this problem, this was a foreign concept to them. They operate with a very different mind-set, that everything is marked to market. They have the ability to close out all of their positions at any time. Their view of the long-term future is the end of the quarter. When we look at a life-

insurance company offering the guarantee on a variable annuity, the horizon is 10 years, 20 years, or more. With hedge programs that involve significant future transactions, you need to understand exactly how these are going to perform under a wide range of environments. We need to understand how they perform within the regulatory and accounting context of a life insurance company.

Just to stress this point, if we think of the GAAP income statement of an insurance company, there are many market-sensitive pieces. Often, VA hedging programs start by focusing on economic risk of guarantees, looking at the guarantees in isolation, valuing them in a market-consistent manner, and assembling a portfolio of hedge assets that can offset that change in economic liability. That is a very challenging exercise in and of itself. But, unfortunately, we need to do much more. If you think of the GAAP income statement, you have a whole range of market-sensitive elements—your fee stream, DAC with unlocking and amortization, GAAP, SOP 3-1 reserves, FAS 133 reserves, and the hedge gains and losses themselves. SOP 3-1 reserves are very difficult. They're not risk-neutral, market-consistent values, but they are stochastic, market-sensitive values. And the market sensitivity varies based on the age of the business and various characteristics of the business.

That sensitivity causes you to make trade-offs between full hedging on an economic basis and hedging your GAAP income. FAS 133 reserves bring in more stochastic elements—namely, interest rates and the whole volatility surface. In addition, people often are looking at guarantee hedging programs now. But then, in structuring the guarantee-hedging programs, they decide how they're going to use them to help mitigate some of the market sensitivity in the DAC process itself. Mungan Slides 5 and 6 show an example of a prospective simulation over a range of scenarios for a diversified block of VA business and the range of quarterly GAAP income on an unhedged basis. There is some volatility in results, and you can have results over a wide range. This happens even with the techniques that are often used in the DAC process to mitigate some of the true market sensitivity that's there on an economic basis.

Next, we take an economic-hedging program and throw it into the mix. You can see that volatility has declined somewhat. But we're still a long way from perfect (if you will), stable and predictable earnings. If we consider the point of view of the CEO or the CFO of the firm, they wonder why we have VA guarantees and VA-guarantee hedging programs in the first place. It's very simple. We use them to increase sales and profits. And those profits need to be stable, growing profits. That's the message that the analyst community wants to get. That's the message that the investors and insurance company want to see—that investing in an insurance company is a fairly conservative investment, a fairly predictable one. That gives us an incentive to meet that goal. If we don't meet that goal, that's going to change the mind-set. It's going to change the valuations of insurance companies themselves.

Mungan Slide 7 shows an example for which we're doing prospective stochastic calculations showing the volatility in economic claims on an unhedged basis across a wide range of scenarios, and then doing that on a hedged basis in Mungan Slide 8. Getting this right involves a whole range of calculations, not just the fair value of your net liability at future points in time. Typically, we would measure these every week for 10- or 20-year scenarios. Then we would do that over a range of stochastic scenarios. But you also need to understand the whole panoply of market-sensitivity measures, delta with respect to all of the major equity indices, rho with respect to short-, medium- and long-term rates. Typically, you'd measure in key-rate rho, then key-rate vega, so that you can understand the sensitivity and the volatility at different points in the curve. Once you have that massive amount of information, you can use that to understand how you can use it to effectively hedge the liability.

Mungan Slide 9 shows a case study with a block of business issued over a range of years with varying in- and out-of-the-moneyness and base product types, and some treatment on the underlying components of GAAP income. If we look at different levels of economic hedging and how they help to stabilize quarterly GAAP income over a range of simple market shocks, this can be the start of the discussion of what type of hedge program you want. If you have full economic hedging, when you factor that into your total GAAP statement for the VA block, you actually reverse the market sensitivity so that your income does better in down markets than up markets. If you have no economic hedging, then you run the risk that the investment community will come to see insurance companies as bets on the market. And as bets on the market, they hardly would deserve high valuations. There is a whole range of intermediate points that give much more stable profiles.

I had many discussions with rating agencies that are very blunt and knowledgeable about hedging of VA guarantees. One senior person at one of the major rating agencies told me that it's cold comfort to see how hedging is going to perform on a computer. They want to see how hedging has performed in real life. If we are looking at any kind of prospective simulations, we want maximum transparency in those kinds of results so that we can know what assumptions went in and how that gets wrapped up in the results that come out. To address the concerns of rating agencies and for senior management to understand how the hedging programs have been working, there's a tremendous amount of effort going into performance analysis to track if the hedge program is doing what it was designed to do and if it costs what it was expected to cost.

Mungan Slide 14 shows a performance-attribution report. For typical weekly or monthly results on a hedge program, you can regard the hedge program as being part of its own profit center, if you will. Often it's a profit center with zero profit margin. Your goal is to have the profit and loss (P&L) come out at zero every single period. That's a fantasy. That would never happen, because there are too many moving pieces, and no hedging program is perfect. But the hope, and, in fact, the historical experience has been that many of these moving pieces represent random

noise that is often uncorrelated, so you do get fairly good results in a range around zero. You want to understand all of the various pieces.

Mungan Slide 15 shows monthly calculations for some programs. The net gains and losses were very small and attributed to their various underlying components. Interestingly, the largest numbers on the table are related to interest rates, not equity. Hedge effectiveness gets more into analysis of the capital-market side of the equation and tracking the performance of specific capital-market instruments relative to their task that they were assigned in the hedge program. In Mungan Slide 15, we show hedge effectiveness versus the three main components of equity-market movements, interest rates and volatility.

I want to say a little bit about the increasing use of options. I think that this is an excellent trend. We've seen this throughout the industry. A lot of banks are much more open today in supplying data and understanding the options market. The insurance companies have taken them up on that, largely using the most liquid, transparent structures in the market. I think that makes perfect sense. The goal is to piece liquid, transparent instruments together in relatively simple, reasonable combinations to have a more-effective hedge of liabilities. It's helped that the pricing in the over-the-counter market has become much more efficient. We've seen transaction costs drop by as much as half. That has had nothing to do at all with life insurance companies. It's been completely in the hands of hedge funds. The amazing increase in hedge-fund activity and the volume of trading that comes along with it have driven down the price for life insurance companies. That's a wonderful side benefit.

Mungan Slide 17 shows another way of looking at performance over time. We look at the net P&L on the hedge program versus the liability that it was hedging. Had you gotten a perfect program, everything would have fallen exactly on a diagonal line. This diagram shows more than a thousand scenarios. This is good for identifying the outlyers. The gray dots are unhedged. There's a whole range of scenarios. You could think of that as your bottom 10 percentile, if you will. Similar to the RBC, you'd see how it performed with and without hedging. The black line (which is the cluster of a thousand hedge dots) is not at all flat, but it does confine your results to a nice band.

Insurance companies have made many strides in increasing the sophistication of their ability to hedge. When I first started doing this back in 1998 and 1999, the state-of-the-art was Black-Scholes-type models, geometric Brownian motion with a single interest rate and a single volatility. If I look at where we are today with term structures of interest rates, entire volatility surfaces, stochastic models that capture all of these factors simultaneously, it's a night-and-day difference. The latest advance of many companies has been a true understanding of the whole volatility surface. If you're going to trade options, you need to understand the price dynamics of those options, or else you're fooling yourself in terms of the economic benefit and costs that you're going to realize in future transactions. The volatility

surface is very similar to all of the accumulated knowledge that actuaries have developed on the yield curve. It's subject to parallel shocks, changes in slope, changes in curvature. If we think of this as a factor analysis, these three factors are orthogonal. We've gone from a simple yield curve into a three-dimensional space as shown in Mungan Slide 19. It's a concept that's familiar to actuaries. It's just a bit more complicated.

To wrap up, I want to make the point that best practices in the industry have advanced very rapidly, and it's been a privilege to be part of it. We've seen much greater sophistication in accuracy in tailoring hedge programs to meet the targets of senior management. This has captured the attention of the CEOs and CFOs at the top of the company. They can start focusing their hedge program very carefully on GAAP income and have greater certainty in achieving specific financial objectives. Performance attribution has been going on for several years. Many of our clients have, for example, three years of performance-attribution results. That establishes credibility and legitimacy in a rating agency that, honestly, prospective-simulation models don't.

Finally, the industry has advanced to using the full range of hedge assets to lower the cost of hedging and increase hedging efficiency. That's been a wonderful advance. What's next? We're going global. Right now, some companies are starting to see VA sales in Asia outstrip VA sales in North America. There's a whole range of products coming through the pipeline for the U.K. and Europe. The leading companies are able to take advantage of their platforms to offer products that are adapted to local markets, but from an economic-risk perspective that are essentially very similar, and then hedge them across all markets with one global platform. That's what's going to occupy us for the next several years. It's going to be very exciting for everyone involved.

FROM THE FLOOR: Ms. Cooper showed a scenario in which the stock market went down 10 percent a year. I'd like to know if part of your hedge was dynamic replication. You had a base, European option, but then you replicated the rest from a delta and rho perspective. In the scenario, how did you address the fact that if the market goes down 10 percent each year in value, that would never happen, and dynamic-replication program will depend heavily on how exactly the market moves on a daily basis. How did you address that?

MS. COOPER: This was intended to be a simplified example. Prospectively, these are intended to be illustrations, and this is an issue. I really showed this example, because, as we're all looking at C-3 Phase II, we're grappling with the fact that we're going to need to address this. The decisions that we make as we dynamically rebalance vary, depending on the changes in the capital markets, the changes in our assumptions, and what's available to trade. I think we need to do the same thing that we do in our ALM models for fixed income. I think we come up with the best sort of algorithm that's appropriate and attempt to look at alternatives. I don't

think that there's an answer. We know that the way that we illustrate what's going to happen tomorrow, a month from now, etc., in our model is, in fact, just a model.