2005 VALUATION ACTUARY SYMPOSIUM*

Orlando, Fla.
September 22–23, 2005

Session 34PD
Capital and Hedge Modeling for Variable Annuities

Moderator: DAVID HOPEWELL

Panelists: DAVID HOPEWELL
HUBERT B. MUELLER
ULRICH STENGELE

Summary: This session discusses capital and hedge modeling for variable annuities. Topics include: What are possible approaches to a capital calculation with hedging? How do different approaches compare? Also discussed will be key issues in developing scenarios and types of hedge strategies that may be employed. At the conclusion of this session, attendees identify issues related to effective modeling of capital when hedging, recognize benefits and drawback of different modeling approaches and recognize major issues related to hedge strategy.

MR. DAVID HOPEWELL: Hubert Mueller is a principal with Tillinghast in the Hartford office. He is the head of the financial risk management effort in North America. He has a Master's degree in probability and statistics from Syracuse. Hubert has an M.S. in applied math from the University of Ulm in Germany.

Ulrich Stengele has a Master's in pure math from Syracuse University and a Master's in applied math from the University of Ulm. I guess that qualifies as impure math in the eyes of mathematicians. He also holds a chartered financial analyst designation and is a recent graduate of Paul Wilmott's Certificate in Quantitative Finance Program. Uli has worked for Nationwide in Columbus for 10 years and has been part of their asset-liability modeling (ALM) group for the last eight. He has most recently been involved in Nationwide's equity hedging risk management activities.

I work for Ernst & Young out of the Philadelphia office. I'm part of their financial risk management practice, at least that's my main area of activity. We are going to

Note: The chart(s) referred to in the text can be found at http://handouts.soa.org/conted/cearchive/valact05/034bk.pdf.

talk about capital, economic capital and hedging and variable annuities (VAs). We're going to spend as little time as possible talking about C3 Phase II specifically and as much time as possible talking about capital as a trend in the industry, as a series of choices. We're also going to touch on things that we can do so that it's not just necessary, but it might even be useful.

We have some learning objectives. We're going to identify some of the major capital initiatives around the world. These things are going on more outside the United States than within the United States right now, but the United States is catching up and certainly will continue to do so over time.

We're going to talk about some of the choices that are involved in constructing capital measures. There are currently very few rules around what economic capital really is and how you calculate it. We're going to talk about some of the decision points that modelers face as they create these models and measurements. I'm going to talk about nested stochastics. How many people in the audience are familiar with the term nested stochastics? (About half or more of the audience responded.) Nested stochastics are very likely to be the dominant technique required to model capital with hedges regardless of the other decisions made. We're going to talk a little bit about aggregation and particularly the requirement to effectively aggregate equity risks with credit risks. We're going to talk a little bit about C3 Phase II, so we knew we couldn't avoid that, but we'll try to keep it in context. We'll spend a lot of time going over some of the decisions involved in designing and implementing a VA hedging program.

I found in a not necessarily comprehensive search, non-capital initiatives. Those of you who actually read it will recognize that BASEL II is not an insurance industry initiative, it's a banking initiative. Some of the regulators, particularly regulators that jointly have responsibility for banks and insurance companies, are asking themselves the question and asking insurance companies they regulate the question: If the banks can do it, why can't you? It's not clear to me that there's a good answer to that other than we just can't do it yet, and that, in fact, is the answer that companies have been giving. They have been with some degree of urgency building out. There's been more urgency in the United Kingdom; there is some now in Canada; there is probably a little less in Europe, but all of the major locations for the insurance industry have significant capital initiatives.

The International Academy of Actuaries is doing some work on C3 Phase II. The U.K. Financial Services Authority (FSA) has published a large document called CP 195 on how life companies need to do Solvency 2 in Europe and the United Kingdom. There's also something called the Prudential Source Book in the United Kingdom. There are Dutch and Swiss statutory capital initiatives. The U.S. statute, SCL2, is just beginning and AUSTI Guideline A2 is another instance of an integrated regulator saying if the banks can do it, why can't you?
I wanted to point out what the world other than the United States is saying about capital initiatives and not only what they mean but how companies are expected to

use them. For instance, on Solvency 2, Fitch, a rating agency, says there is a risk that those who have not taken any action so far are left totally unprepared. Their view is, and this is probably true of the other rating agencies, that companies that do not have risk measurement systems are sophisticated and realistic and are used for management decision making. That's a very important point. These are not compliance exercises; these are intended to be management decision-making exercises. They are weaker from the perspective of risk management than those companies that do. The more sophisticated models that are adequately integrated into business management will result in us giving individual capital guidance which is lower than would otherwise be the case. We're not talking about compliance here; we're talking about money. If you do it better, you don't have to hold as much capital. That's the reason to do it.

Unlike the banking sector, the insurance industry lacks an international standard for assessing solvency. AUSTI goes on to say that they have a preferred approach. If you want to read about it, go to the AUSTI Web site. It also says that it doesn't expect insurance companies to be able to do it and so it gives an intermediate stage approach which it will consider acceptable for a while. That's probably not a road map, but at least some guidance on the way this might play out over time in North America.

I'd like to talk about economic capital. How do you do it? Why do you do it? What's the point? Is it a compliance exercise? Does it measure something that management cares about and can management use it effectively to make decisions and to adjust the risk stance and the risk-adjusted return expectations of the company? There are a couple of ways to approach economic capital. The first is: Are you measuring the solvency of the company or are you measuring the viability of the company as a going concern? Solvency is sort of a closed block concept that regulators like. Are there enough assets to run off the existing business without further injections of money? Going concern is a somewhat stricter requirement actually because as we know, most insurance companies are not going concerns once they get to BBB credit ratings. You can get to BBB and still be solvent.

There are issues of time frame and these issues tie into both what you think the purpose of the measurement is and how you intend to go about calculating it. There's some feedback between the decision on what systems are we going to use to calculate and then how do we choose our measurement approach and what is the purpose of the calculation. Are we interested once again in run off, that is surplus efficiency over life time, or are we interested in what's our balance sheet going to look like in a year when the rating agencies or others who have a say, look at it? What's the likelihood that we will still have positive capital a year from now?

Once again this goes back to the purpose. If you're a regulator, you care about the long-term statutory balance sheet. If you're a company, you may very well care about the GAAP balance sheet. If you're in Europe and you're doing market-consistent ED, you're thinking in more of a market consistent space. Are your

options and guarantees valued at something that's an estimate of their manufacturing value? Finally, if you're a company on a going concern basis, you might choose to think about what your forward appraisal value is, which is somewhere in between the statutory fully funded run off and market-consistent measurement. In a year, will I be able to sell this block of business for break even or better? Obviously the valuation at which that happens is different from the others, and it's also arguably a market consistent valuation for a block rather than a set of options.

One of the main issues in economic capital from a company perspective is diversification. There are a lot of decisions that need to be made in terms of how product lines and asset classes are going to diversify as you roll up from a single business line into a whole company. There's also the question of how hedging is reflected, not because hedging is so hard, although after you listen to Uli, you will probably have respect for his skills, but because it requires the nested stochastic approach with a set of market-consistent assumptions rolled forward over time, which is a fairly difficult thing to construct.

Then there is the question of capital. Once you make those decisions and decide, the question is: What is enough? There have been several large draw downs in the U.S. stock market. Draw downs top to bottom of almost 50 percent in the last 100 years. That tells me that if you have a scenario set that goes out 30 years, about half of those scenario sets should have a top to bottom run of almost 50 percent. That's not generally something which comes out of the standard lognormal fixed volatility kind of scenario generation approach, and in fact, it's this top to bottom movement that tends to be one of the main drivers of risk for a number of reasons. The first is that the top sets the expectation in terms of income. The top sets the level of guarantees. The top also encourages more risk taking, because on a backward-looking basis risk taking has paid off so well, and if you haven't done it, you've fallen behind.

The problem with this is, that top to bottom it is a conditional tail expectation (CTE) measure. I didn't say C3 Phase II, but this is how I calculated it. It can be, depending on the benefit structure, three to 10 times what it is at the top. If, at the top of a bull market, you calculate 2 percent as your required capital, at the bottom you may look at 10 percent, 15 percent or 20 percent on a reduced income base. Funding that increase in capital is a significant issue for companies and that's what's driving them to hedging, although that's not the only way to fund it. You have to ask yourself what can be a product margin to support over product life. Are there other product lines that are truly uncorrelated with the capital requirements, the VAs and the income characteristics of VAs that can fund the increasing capital out of their profits? At what point is it impossible to rely on diversification and capital? If you don't have enough diversification, capital by itself is not a sufficient mitigant for these risks. Products don't have the margins to support the maximum capital over a business cycle, and by business cycle, I mean the long cycle that leads to a large bear market. In fact, the FSA in the United Kingdom says explicitly

that if you're not diversified, we don't care how much capital you hold, it's inadequate, you have to do something else. The two main tools for mitigation are reinsurance and hedging.

Having said that, it's clear that if we're subjected to these capital rules, this will simultaneously drive management's pricing desires, risk appetite, strategies and their hope for diversification. Then the question is: Once you know that this is an outcome of the capital measurement, can you take a required capital measurement and turn it into something useful so that you can actually make actionable decisions along these dimensions of risk appetite, diversification and pricing? Does that allow you then to make hedging or mitigation decisions and to test the cost effectiveness of those mitigation decisions against each other? That's what I would consider useful.

In order to get there, there are some tough things you have to do that typically compliance exercises have not had to do. You have to have management metrics. You have to do it quickly enough that you could get a rate over a decision-making cycle, because you're not going to get the answer that you want the first time. There are too many variables. You need consistency, that is you need your metrics to point generally in the same direction, if not give you the same answer, which they won't.

Derivatives in reinsurance are traditionally considered risk mitigation. As an ex-derivatives guy, I think of them as diversification. What you're doing is you're swapping stocks for bonds or risky assets for bonds, so derivatives are just a way to control your equity to fixed income. Once you do these things, how do you prove it to the authorities? The authorities might be a regulator, an activist regulator or a rating agency. They're interested in this. They're looking for companies to build their own models, make decisions based on these models and then to prove to them that the models are effective and that the decisions were made in a way that's consistent with the output of the model. That's considered to be a good enterprise risk management process or at least a necessary part of a good enterprise risk management process.

We know that for VAs there isn't much reinsurance right now. There probably will be some. It's not going to look like the reinsurance of old. It's going to look more like options, and I think that's a fair bet. It's going to be priced more like options and so it's going to have to be assessed very much in the same way a hedging strategy would. Reinsurance will look more like another form of hedging, with the same demands on infrastructure, with the same demands on modeling, with the same requirement for understanding of what the risk transfer actually is from a capital markets perspective, as well as a mortality and morbidity perspective.

That leads up to modeling. Hubert and Uli are both going to talk a lot about modeling, so I don't want to beat this up too much, but it's probably the most important takeaway from this meeting that high quality modeling is required both

to calculate these numbers and also to create credibility among external constituencies. That's a pretty critical outcome, because this is an awful lot of work to do if you don't get any credit for it. One of the main benefits again is a dollar of capital in the past has been a dollar of capital more or less. In the future it probably won't be. Capital held without sufficient institutional controls and modeling will be haircut. A dollar won't be a dollar any more.

I work with significant compromises that are required to get these kinds of answers. They're either slow or you have to work around or you have to leave out the derivatives and calculate them to the side and so on. You can get answers that way. You can probably do a compliance exercise that way. I'm not sure that you'll be able to continue to do the rating agency and external constituent kinds of work that will need to be done in the future unless you can prove somehow that your approximations are close enough to correct that they're okay. I don't know how you prove that without actually doing the modeling.

One of the main challenges is the nested stochastic structure, along with the requirement that comes along with the nested stochastic structure of switching between assumption sets. You have nested stochastics that go with real world assumptions for the base, but the nests are risk neutral. You have risk-neutral within risk neutral. You'll have C3, which is a kind of a real world assumption set and then risk neutral inside it. There are a number of mixes of scenarios. You need to consider not only mechanically producing those, but creating rational transitions when you go from one world. The host scenario to the risk-neutral scenarios is a bit of a challenge. Then finally traditional actuarial grouping is going to be under pressure. I'm not sure it will go away entirely, but the dimensionality of VAs with guarantees is so high that in general it will be quite difficult to come up with very effective groupings that significantly reduce your calculations.

There are three hedging approaches I would say. I'm going to call one static hedging and Greeks matching. For those of you who have seen other terminology, static hedging is one period hedging. That is, you start at the beginning and you end at some point well into the future and the hedge stays in place and roughly unchanged during that time. It tends to be a little bit less accurate on a model results basis, because you're not adjusting these things from period to period. On the other hand, it's much more robust against the assumptions of rolling over the hedge, which tends to be a large and hard-to-get-your-arms-around source of risk in these programs. Frequently it's wished away by assuming sort of static capital market assumptions. They clearly are not static. I've built a hedge for five years; it's not static, trust me.

The static hedging is also something that can be done actually in many cases, without resorting to a separate Greeks calculation. You treat it like a portfolio management construction and try to minimize some error function between your hedged target and your option portfolio. That's pretty easy to do. You can do small examples on a spreadsheet and it once again doesn't require you to maintain a

separate set of risk-neutral assumptions, to generate a whole other set of scenarios and to gain the knowledge required to generate highly realistic risk-neutral scenarios, which is actually a fairly hard thing to do.

Greeks matching will probably continue to be the majority or the main approach for hedging. The issue with Greeks matching is there is not any kind of resemblance between the capital measure and the risk-neutral measure in terms of the riskiness of the guarantee or an option. A Greeks-based hedge is probably going to be inefficient from a cost-benefit perspective especially if that cost-benefit perspective includes the cost of capital.

The third problem with Greeks matching (it will have to be done and this problem will have a lot of resources and effort thrown at it) is the iterative calculations in nested stochastics. It's more than just boxes with CPUs, but that's a big part of it.

You have to generate scenarios. You can find tools that will generate real world or C3 Phase II or risk neutral or you can build your own. In general, I'm not aware of any tools that will handle the transitions between one assumption set and another, particularly on a real world risk-neutral basis. That's going to become a bigger issue over time. Of course, we know tail measures are really sensitive to the assumptions in your scenario set. That's one of the issues. For instance, a tail measure is very sensitive to that draw down from the top of the stock market to the bottom of the stock market, because that's what drives your tail particularly if like most companies while the stock market's going up, you dividend out your earnings, where are they when the market goes down, how do you get them? They're either spent or they're somewhere else.

I think the last bear market taught us that equities and credit risky bonds can be a lot alike under stress. There have been seven bear markets. There have been seven of those events since 1960 where the market has dropped more than 20 percent from its recent high. In five of them, the AA spreads increased by more than 100 basis points. Those are roughly simultaneous spikes of credit spreads and equity. This is a fact. Any model that doesn't incorporate this on an integrated basis is not a realistic model. I think there are both empirical and theoretical arguments as to why this has to be.

The problem with it is, if you put this stuff on a spreadsheet and do correlation, you get a number very close to zero. The reason for that is they're only correlated when things go wrong. Credit and equity are 60 to 80 percent correlated when things go wrong and 0 percent when things are going well. Don't be fooled.

MR. HUBERT B. MUELLER: First of all, how are companies defining economic capital? Probably the most common methodology for defining economic capital is to look at the market value of assets minus liabilities and to look at the probability of ruin. You have a probability of ruin without economic capital and then you say: What is my target that I can afford? Is it rating agency driven or is it regulatory

driven or otherwise company driven? You basically define your economic capital such that your probability of ruin is below that target. If you're setting yourself a target of 1 percent, maybe because the rating agency is measuring you at the 99th percentile, well that would determine your economic capital for that line of business or for the company in total. Of course, if you're modeling a line of business at a time, you're not incorporating the benefits of diversification; you're modeling all lines of business together and you are essentially integrating and incorporating diversification. That goes past the hurdle of having these correlation matrixes. Some approaches in the past have tended to use correlation matrixes, which is a huge nightmare, because you have a number of risks, let's say 500 different risks going across and now think about how many elements you have to fill out in that correlation matrix and not just once a year, but also at the beginning and end of the year. When I have seen those, there were a lot of zeroes and ones, which surely doesn't give a lot of information. I think it's worth modeling this and looking at diversification in benefits by actually modeling on an integrated basis.

There are a variety of approaches and uses for companies modeling and measuring economic capital. The most common ones include economic scenarios, modeling the stock and bond markets for the last 20 years. It gives you a pretty good idea, because we've seen quite a bit of volatility going all the way back to 1987 and then the first couple years of this decade. Some companies and also some regulators have specified stress testing rules. In Germany, for example, companies have had quite a lot of volatility in their capital because of the equity market downturn. They were much more invested in equities than the U.S. companies have been, not just for capital, but also for backing liabilities. As a result, the German regulators and actuarial association have put a stress test in place which companies have to fulfill. It's a one-year stress test and the assumption is the market goes down by 35 percent and interest rates increase by 2 percent. Companies have to show that they follow under that scenario over the next year. That's been pretty stringent in the past for companies to fulfill. Quite a few companies had to shore up their capital to qualify.

Some companies use FACTA tables, although I wouldn't necessarily recommend that. The most common approach certainly is to use stochastic models, using both risk-free and real world assumptions. We're going to get to more of that, so I'm not going to spend too much time now. Also, if you're thinking of a world with European options, single exercise, you can use option pricing methods, mean variance and co-variance models. David talked about credit risk methods. Increasingly we're also noticing that the analysis includes not just financial risk, but also operational risks which are becoming much more common. In fact, if you look at the failures in the market, both in the United States and other markets, a lot of them were driven by operational risks, not by financial risks. Leaving them out is certainly a shortcoming in your approach.

Typically, as I said, economic capital is based on the probability of ruin. Keeping ruin to the probability below a certain target level, that tends to also now become the preferred method by the rating agencies, and I'll get to that in a second. What's

worth mentioning too is when you define economic capital based on shortfall risk, like either what we just talked about or CTE-type measures, it's also a coherent risk measure. Coherent means if you combine risks, then the risk from combining those events is actually less than or equal to the measure of each risk separately. Now value at risk (VAR), which is what the banks use for defining economic capital, does not fulfill that property. It's not a coherent risk measure, and the discussion has gone both ways. Banks are claiming that if you combine risks it can actually be worse than each one of them separately, whereas I think the insurance world is probably more the opposite and says if you combine risks it should be better than looking at each one of them separately. I think the truth is not necessarily going one way or the other.

As I mentioned before, the regulators and rating agencies certainly have spent much more attention on defining appropriate measures for capital. Certainly when companies use economic capital measures, they tend to be much more proprietary. What we see in the industry is the larger companies are developing proprietary capital models, because the existing risk-based capital (RBC) formulas tend to be industry specific, so it's an average over all companies. What you would expect is the larger companies tend to be a little better than average. The smaller companies tend to be a little below the average. The larger companies have more incentive to build proprietary models and have certainly driven the charge there. They also tend to be prospective methodology. Now, the regulators traditionally have tended to use industry formulas, not company specific, purely formulaic, and purely retrospective. It's kind of a lack of reacting to market events.

Of course, we all know that principles-based regulation is here. This meeting is all about it. One of the reasons for that is when the RBC formulas being used in the industry were developed, a lot of the risks that you currently have as products were not available. They were not around. Fifteen years ago when a lot of the RBC formulas were developed, there were no living benefits on VAs. Today almost all the VA products have living benefit guarantees. It's also worth mentioning beyond what David was saying that there's a significant effort underway in Europe to sort of take what Solvency 2 specifies at a high level and make it more pragmatic as to what companies' capital models have to look like for them to be admissible to the regulators as part of Solvency 2. There's a group of about 13 European large multinationals where the chief risk officers have gotten together and established guidelines for calculation of economic capital and diversification risk and how they would be acceptable. Their focus is to try to develop some common guidelines that the regulators in all those European markets would accept. I believe their Web site is called CROForum.org for those of you who are interested. One of the publications called "Principles for Regulatory Admissibility of Internal Models" was published this year. I would highly recommend if you are working in this area that you take a look at that. It's very interesting.

Of course, the reconciliation might be coming to the rating agencies as well, because the rating agencies have quickly caught on to the developments in the

market and are now looking at going away from the industry-based formula methods that they've been using in the past for capital adequacy and looking at company-specific models. Standard & Poor's (S&P) started with a financial products model and is now actually going beyond and saying we're going to take that approach from C3 Phase II and we're going to look at proprietary models from each line of business. Show that this is really a sound model. Specify the model, the assumptions and the methodologies. Look at using percentages of a capital adequacy ratio, which some companies have been doing in the past for certain ratings. They're now looking at the tail risk and the percentile that they're looking at is consistent with the company's desires or current financial strength rating. For example, at the spring meeting we had somebody from S&P talk about how the financial strength rating of AA would translate into the 99.5th percentile. The way they rationalized it was to say roughly a AA bond has 10 basis points annual default risk. If the duration of the business is five years, you take .999 to the 5th or simply multiply the 10 basis points times 5 and you get to 99.5. That's one way to do it, because we all know the C3 Phase II formulas. The 90 CTE level is for 100 percent RBC; it's not for a AA or a A or AAA rating, it's a BBB equivalent. Moody's is also coming out with a new model. They've already introduced the property/casualty model and the life model's expected to be out this year. It will also be a stochastic model. Fitch is working on one too and has actually just published something for VAs. For those of you who are interested you can go to the Fitch Web site. They're using CTE levels with higher CTE levels for better ratings, and they are exclusively talking about giving hedging credit. There's a lot going on.

Let's look at how you would reflect capital and hedging and placing. There are lots of reasons why you have to do this today, not just guaranteed minimum death benefits (GMDBs) and guaranteed living benefits (GLBs). There are secondary guarantees. We all know about the UL working group development, the rating agencies are doing it. Also if you are looking at developing capital and pricing, it might require you not just to do that at issue, but also at future points in time. You are talking about a real world nested stochastic projection. If you're going to reflect hedging in pricing, then you might be talking about what David called a double nested stochastic, which is real world and risk neutral. Of course, it gets pretty complicated and I know some of you are looking at your computing capacities and saying I'm not going to be around when the results arrive. Well that's unfortunate.

In theory, you want to do all these things perfectly and have lots of model points and lots of scenarios, but if you want to do that, you're looking at run times and results which are going to be out way past when you want to see them. We had a pricing size recently where what the company wanted us to do would have taken four years to run on a bank of four computers. That's not really acceptable. By that time the market's gone. You have to look at methods of building representative models of a compressed number of model points just to look at ways to reduce your scenarios. There are lots of ways to do that if you know what you're trying to get.

If you're developing a financial projection and you want to get a best estimate of future profits, you can look at the mean of results. You don't necessarily have to look at the tail risks. If you're looking to develop capital or you're trying to hedge risks, increasingly you have to focus on the tail, so you have to figure out which scenarios are giving you the tail. We've done some analysis where you can actually replicate the scenarios, 1,000 scenarios that are supposedly CTE 90 with one scenario. If you do it for a very small subset, you can figure out what's giving you the tail and you only have to run that one scenario. There are lots of ways to really improve your run times doing that. Of course, the hardware is part of it too. Model design is part of it. There are lots of ways the companies are looking at designing, developing and improving hardware. Distributive processing is becoming very common. We know one company published at Hartford Life talked about it publicly using 350 CPUs, using good computing. One company at the meeting told me they're using 500 using good computing. These are numbers that probably a year or two ago nobody would have fathomed and I don't even think the banks use numbers like that, because they don't have that complexity.

There are other levers too, like the number of model points, projection periods, even things like how often you calculate future events. For example, you can model a hedging strategy by doing a weekly simulation of a risk-neutral projection, real world on risk neutral. You can also do it monthly or quarterly or annually. Of course, the less you do it, the less your accuracy. You have to find the right balance between the two and you have to revisit that model optimization throughout the model-building process.

Now, when you talk about stochastic pricing, we're digging one level deeper, some of the things you have to think about are dynamic assumptions. It doesn't make sense to use stochastic scenarios and have static assumptions. It is like the saying: "Garbage in, garbage out." It doesn't mean anything if you do that. The market experience on that is still emerging and I'm going to talk about that a little bit more. The scenarios that you're developing should be integrated. It doesn't make sense to use interest rate and equity scenarios which are not correlated and then do hedging runs. You're going to get abnormal results. The risk-neutral and real world scenarios have to be consistent, not just correlation between the equities and bond rates, but they also have to be consistent. If your outer loop is real world and your inner world is risk neutral, then whatever point you go from real world to risk neutral you have to have consistent volatility and swap rates, otherwise the results don't make any sense. One company came to us and said, "Well, we're getting higher results after hedging than before." I thought either you're doing a really poor job at hedging or something's wrong with your scenarios, and, of course, the scenarios were not consistent.

When you look at how you define capital, you basically have to do the runs in several iterations now. You first do an analysis of the tail risk of the product with no capital. You figure out the tail risk and then you include capital based on that tail risk. You have to do several iterations and there are various ways, of course, to

define capital. You look at your results. You have to declare on what you are most interested now. Are you most interested in the mean results or the 80th or the 85th percentile or the very end of the tail, the 99th percentile? That will also drive the number of scenarios. It doesn't make sense to use 200 scenarios if you're looking at a 99th percentile tail risk. You don't have to do a lot of math and a lot of statistics to figure that out, that it's not statistically confident. You do that five times and you get different results five times. You do 1,000 a couple of times and you're going to get pretty stable results even at the 99th or 95th percentile. For 99th I would almost recommend to do several thousand scenarios if you want to get stable results.

Of course there's an increasing need to incorporate hedging. Why do we need dynamic assumptions? If you haven't been to any sessions where this has been discussed, here is a quick synopsis. First of all, dynamic assumptions are used for the actual utilization of living benefits. If you want to model a guaranteed minimum withdrawal benefit (GMWB), you can make the utilization of that benefit contingent on several factors—for example, how much is the guarantee on the money or what was the market return over the past year? What you would expect logically is that the more the guarantee goes into the money, the more people would utilize the benefits. That's probably one of the cautions with the current living benefit modeling. Don't purely make your assumptions for living benefits consistent with what you've seen the last two years. We don't have a lot of record and history with this yet. The last two years the market has been pretty good, so you would expect to see much higher utilization of living benefits as we go forward if the market goes south or at least is stagnant for a while. Since introduction of living benefits, we really haven't seen very adverse markets yet. I would caution companies to think utilization is 25 percent, let's use that going forward. We've done stimulations where we've done all the way up to 100 percent. You can segment your policy holders and say you're going to segment them into the people that didn't buy the policy for that benefit and just kind of got it sold, and the people that really bought it, the income buyers that are going to be a much higher utilization regardless of where the market goes. You can specify it that way.

There are also things like dynamic lapses and partial withdrawals, not just dollar for dollar withdrawals, but in general partial withdrawals can be dynamic, or the utilization of resets. A lot of the living benefits today have reset features every five or seven years. You can make that contingent on how much the guarantee is in the money or where the market is relative to where it was or where the account value is relative to where it was when the policy was bought. Of course the industry experience is still emerging. Sensitivity testing is key. Don't just use one set of assumptions, use multiple. Try a few. You're going to see some dependence of the results from the assumptions, and there are some studies currently underway. One is by the Society. The Risk Management Section is doing a survey on policyholder behavior in the tail section. LIMRA is doing one too.

When you include capital in the pricing, the traditional approach used to be to use something like 200 or 300 percent RBC. It used to be 200. Then it went to 250.

More recently it was probably more like 300. The larger companies now increasingly use a percentile-based approach consistent with their financial strength rating, so the higher your desired rating that you have or that you want to maintain, the further out you want to go on the tail for a loss distribution to figure out how much capital you need, so the more capital is costing you.

If you want to do this step by step, at issue you would run, let's say, 1,000 real world scenarios. You run your base product without the benefit and let's say you're looking at the 99.5th percentile, consistent with the S&P AA rating that I just showed you. Then you do the same thing over again where you run the base policy now with the rider benefit and you look at the difference in the loss distribution between the two at that percentile level. Of course, the percentile level should be something where you're statistically confident that this is reasonable. If you're only running 200, I wouldn't suggest doing that, you want to run enough that you get stable results. You then define your capital, which is the difference between the two at that percentile level—for example, basis points over account value. Let's say it's 50 basis points over account value and you put that back into your base run with the rider and now you have a run including the cost of capital. Of course, you could do this at future valuation dates too. I was only talking about times zero at issue. If you do it at future valuation date in the same run, then you're really talking about a stochastic on stochastic calculation using real world just for the capital. If on top of that you want to reflect hedging, then you're talking about a nested stochastic on stochastic calculation where your hedging impact is coming in using risk-neutral scenarios. It gets even more complicated if you have annual premium products. It's simpler with single premium products like VAs. It gets more complicated with variable products. We're now seeing variable products that have living benefit features like GMDBs, guaranteed minimum withdrawal benefits (GMWBs) or guaranteed minimum accumulation benefits (GMABs). It gets more complicated if you want to do that when you have annual premium type features. There's a lot to think about.

An example would be a GMWB model where we looked at the tail risk with no hedging and then various levels of hedging simplified of course. It's delta hedging, then we combine delta and vega and then delta, vega and rho. This was a case study that we did for one company. The hedging improves the tail, but you lose profit, there's no free lunch, hedging comes at a cost. On a present value basis your profits will get lower with increased intensity of hedging, however, since your capital is defined based on the tail risk, you might actually from a profitability perspective do better because you have less capital. Clearly capital is a strain on your results. If you're pricing to achieve a 13 percent return and you're earning 3 percent after tax on capital, every dollar of capital is costing you 10 cents each year, the difference between your target return and your asset tax earned rate. The more capital you need, the more of a drag on returns you get. In fact, in some of the work that we've done, even if you put in a hedging strategy to the level of 30, 40, 50 basis points, if it helps you a lot in reducing the capital, you might actually, from a profitability perspective, come out fairly even because of the reduction in

capital that you get out of it. In this case, after we include the capital, the tail risk actually was, of course, worse with no hedging.

The issue when we talk about hedging and incorporating hedging, as you might have gathered from my comments, is it gets pretty complicated and pretty run-time-intensive. There are ways to simplify what you have to do. First of all, there are ways to approximate the scenarios which provide the tail risks. There are lots of good methods out there. If you're looking for some information there was a stochastic modeling symposium two years ago where there were several papers on representative scenarios. It's a very good methodology where you can group 1,000 scenarios into say 50 or 20 which represent the full 1,000, or in extreme cases even 1 or 2 representing the tail maybe with simplification. You can also simplify the hedging impact. Rather than doing that stochastic on stochastic on stochastic, try to figure out from a smaller subset how much is that hedging costing you, what benefit offset do you get and maybe include an inefficiency factor. We can actually model all that in a base run without having to do nested stochastic. You can use representative scenarios, and also how often you simulate that hedging strategy has clear impact on your run times. Whether you do it weekly or monthly or quarterly or annually have the appropriate multiples. For one model point, if you wanted to do it I would say almost right, near perfect, you'd have a trillion calculations if you wanted to do 1,000 scenarios. Now, I'm not even going to tell you how many hours that would be on a PC, but let's say if you did 1,000 scenarios real world, then at future times capital and then trying to get the hedging impact, that's a billion scenarios and then if you did that for 20 years and 52 weeks because you're doing weekly hedging, you're talking about a trillion calculations. I think that's a little hard to imagine how you would get results in a reasonable time.

Then you have to go back and say: How many scenarios do we really need? Let's figure out what scenarios are giving us the tail and maybe do monthly hedging instead of weekly hedging. When you're down to a billion and a half, we're talking more reasonable numbers now. If you do that even further and figure out what scenarios are giving you the tail, maybe only doing one on the tail and doing 100 on risk neutral, as long as you get reasonable results, you have run times where now you're talking only a day or two on one PC. If you have a band of PCs you can cut the run time by an appropriate factor. It's worth thinking about how you set up your models before you go ahead and do all the busy work and then wait for the results.

I wanted to talk a little bit about C3 Phase II and how capital and hedging plays into that. I'll try not to duplicate things that have already been said at this meeting. I'm assuming you've already gotten some of that in one or the other sessions that you've attended. Clearly we're talking about a pretty sophisticated projection model that you have to have in place. You can't just do something that maybe was good enough for C3 Phase I. It's not good enough any more. You have to understand your product features. You have to understand the CTE 90 methodology. You also have to do a standard scenario. The capital will be based on the higher of the two.

You can get hedging credit, but if you look at the description of how you get hedging credit, it's not quite simple. It's about a 10- or 15-page appendix that describes how a clearly defined hedging strategy can qualify.

If you want to avoid that, there is actually an alternative methodology but you can only use it if you only have death benefits and you will not get hedging credit, which makes a lot of sense. If you can't model, you can't get hedging credit, because you can't really determine how much benefit you get from hedging.

The methodology works like this. It's a quick one. You have a total asset requirement which develops stochastically. You compare it to the model and you take the greater of the two. There are some smoothing and transition rules. You deduct your reserves and you also have a tax adjustment that you have to work through. The CTE 90 is a similar picture that we just looked at before. It's essentially the economic capital equivalent using the average of the worst 10 percent including positive results. There seemed to be some confusion among the audience in one of the other sessions. If you have results on the worst 10 percent which are positive, you don't have to eliminate them to zero. At one point we were talking about a modified CTE 90 which would have only used the tail, the negative results. You can use the positive results if you have any positive ones in the worst 10 percent.

For calculating your total asset requirement including hedging, here is essentially the methodology how you would do that. You calculate your total asset requirement using your best efforts for hedging, using your strategy if you have one in place, assuming if it qualifies as a clearly defined hedging strategy. Step number two is to recalculate that reflecting the impact of risks which you might not completely hedge, or in other ways reflect some inefficiencies, because your first one is most likely going to be subjective. The actual reported total asset requirement is a weighting of the two where you start from best efforts and then you add the difference between your total asset requirement based on the adjusted results and the best effort results. Those of you who printed out the hard copy from the Web site, the formula was reversed, so please correct that. It was adjusted and best efforts were reversed, and this so called inefficiency factor has to be at least 5 percent but could go all the way to 100 percent. That's, of course, a big judgment call for the valuation actuary.

Probably the most important thing is the actuary must provide certification that the assumptions used are reasonable. If you want to use a hedging strategy, it essentially has to be in place for at least three months or you have to use significant back testing to show how it works. For the standard scenario I think it has to be in place for six months, not just three months. Even developing something now might be a little late if you want to use it at year end. I think some companies are over optimistic in how much hedging credit they're expecting to get. If you're expecting to get more than 50 percent hedging credit, I think you're kidding yourself, because if you actually do the analysis and you figure in the cost

of the hedging and some of the inefficiencies, you're probably going to more come out with somewhere in the 25 to 50 percent as a reduction of capital on a net basis everything all in, from all the results that we've seen and from what I heard some companies talk about. In pricing, I've seen some companies use as much as 80 percent. I think you're kidding yourself if you do that.

Uli's going to give us some more detail on that. Incorporating capital and hedging is one-time intensive and quite complex. It does require a stochastic on stochastic model, contingent on tail risk, but there are ways to incorporate the analysis and duty analysis without losing too much accuracy. We think that implementing these models will actually further bifurcate the industry. We're already seeing some of that. The larger companies tend to be more sophisticated. They can now do things on the product side, which maybe some of the smaller companies that don't hedge can't really do any more. What you see is the larger companies have a larger market share and in fact some of the smaller companies are already throwing blocks on the market. There's been quite a bit of mergers and acquisitions (M&A) activity. Some blocks have already changed hands. We're expecting more of that to come, so I would expect some consolidation in the VA market even before the end of 2005.

MR. ULRICH STENGELE: The world's moving to computing capital on a stochastic basis, which frustrates a whole bunch of challenges. Everybody's embracing it, so it's not just the nice thing to do any more. We really have to do it. I'm going to narrow the focus a little bit now and talk about the action between capital calculations and hedging. The two main things that I'll talk about are what I am hedging and how am I hedging it. The main questions are: How do I simulate the hedge strategy and how do I figure out how effective the strategy is? It presumably isn't just important for computing capital, but we really want to know how effective we are in hedging benefits. I'm going to talk about that in terms of two extreme strategies, the static and dynamic strategy. David talked about that a little bit already, so I'm going to fly over that.

What's being hedged? It's probably not as obvious as one might think. I'll give a few examples of what I might hedge. The first one I could hedge is a gross rider cash flow and it can be related to the withdrawal benefit, the death benefit. I could also hedge that net of fees, which some companies are doing. The hedging statutory financial, I can theoretically hedge C3 Phase II capital and my hedge would just simply offset any capital movement in theory. I can also hedge GAAP financials and I think a fair amount of companies are doing that, which obviously is a lot easier if the benefit that I'm hedging falls under FAS 133 rather than the SOP 03-01.

I'm going to think of the hedge strategy in terms of two extremes. One is a buy-and-hold strategy which usually people employ to cover an expected cash flow within a known maturity, and the accumulation benefit is a good example here. The other extreme is dynamic replication. I think of that as going in the basement and

manufacturing an option in the basement rather than buying it from Wall Street. It's dangerous. The dynamic replication is a lot more flexible. I could do that to cover cash flow or I could do that to hedge statutory GAAP financials. What most people are doing is actually a combination of the two, and we call those semi-static strategies.

The central issue is I simulate a hedge. Capital calculations are based on what? They're based on projections of future asset liability cash flows. Question number one is: How do I compute these future cash flows? Question number two is the effectiveness question: How well do the future hedged cash flows achieve the objective of the hedge strategy that I put in place? You've heard both David and Hubert talk about this. It's very complex. I mean at a minimum you have stochastic in stochastic if I just do a capital calculation. If I do pricing, I have three levels of stochastics.

Suppose I have an accumulation benefit. Ten years from now I promise the policyholder if your account value is less than your premium, I'm going to make you whole. Suppose the policyholder is invested in an S&P 500 fund, and finally also suppose that I know this policyholder is going to be around 10 years from now. It's a very simple set-up. In the realistic world for realistic scenarios, in some statistical sense, feel like history, so if I have an equity scenario I look at the mean. If I go back in history I can see that. It's an 8 to 10 percent mean. It's something that we're very familiar with and very comfortable with. In the past when we've priced, we've used realistic scenarios. I can take this accumulation benefit and I can look at 1,000 different S&P equity indexes after 10 years and I can figure out for each of those scenarios what I'm going to look for is how much do I have to pay my policyholder. What I'll find is that for most scenarios actually you don't have to pay anything. Some scenarios I might have to pay a lot. If the market started at 1,000 and said 500 at the end of 10 years, I have to pay the policyholder a lot of money.

The key here is that the outcome isn't certain. There's a whole distribution of results, and I don't know what it is. What we typically have done as actuaries is looked at the mean of that distribution. We probably looked at the tail of the distribution to get comfortable with the risk that we wrote. That was realistic.

Now I'm moving to risk neutral and the set-up for that is I can also go to Wall Street and say give me an option, a put option in this case that pays me exactly what I have to pay to the policyholder at the end of 10 years. Wall Street is probably going to happily do that, and they're going to charge you a price where you'll think that's way higher than what my average realistic scenario was. Well, it should be, because if Wall Street tells you this is what I'm charging you, it's a certain price. There's no distribution associated with it. That's what we're trying to do in hedging. We try to take the fact that there's a distribution and make it tighter in some sense, reduce variance if you will. That's where we use risk-neutral scenarios. Now, there's a lot of math behind it that obviously we're not going to get into here, but I find that a somewhat helpful way of thinking about it.

I start out at time zero. I have to figure out what kind of a hedge I need to put on. I do that by running risk-neutral scenarios. Then I have four different futures at time one. Again I have to update my hedge position, so I'm getting risk-neutral scenarios within realistic scenarios. Hedging risk-neutral C3 Phase II is realistic.

How do I compute future hedge cash flows? They're defined by the hedge strategy and again I have the two extremes: the buy-and-hold strategy where I buy a hedge usually from Wall Street to diffuse an expected liability cash flow within the normal maturity. Group accumulation benefit, the one that I was describing in my risk-neutral realistic example, could be an equity-indexed annuity. Its benefit is I know what the maturity is. The dynamic strategy is more complicated. It involves active trading. It's a lot more flexible. There are the semi-static strategies, which are the combination of the two.

Now I'm going to talk about how I can simulate hedge cash flows to both the static case and the dynamic case. I'm going to start with the static case, which is the easier case. Modeling the hedge in the static case means that I have to mark or it means I have to compute the mark to market at each future model time point. At time zero I already should have that value, because the investment bank from which I bought my static hedge tells me how much it's worth and then the mark today and at every future point in time gives me both income and balance sheet items. We might be lucky in a sense that the static hedge will actually allow me to use a fast form solution for calculating the mark. If I bought more exotic options, that might be a barrier option, a knock in or a known out option. I might have to actually run scenarios which are slower, or I might have a reasonable approximation. The terminology might be confusing again. When I say Monte Carlo valuation, I mean I have to run risk-neutral scenarios to value a hedge, to value an option. Anyway, no matter what valuation method I use, there are always issues even if I have a form method.

As an example, suppose we have a buy-and-hold strategy again. We hold European puts or calls, it doesn't matter which and I can use Black-Scholes to value puts and calls. Black-Scholes needs fixed inputs. The strike is fixed and the maturity date's fixed. So I'm going to come from the economic scenario, the C3 Phase II scenario if I do capital calculations, like the equity index and the interest rates. Some I'm going to have to think about a little more. Think about the dividend rates which are actually quite important when you buy long-term options as well as the implied volatility, which is the most complicated assumption. Just as an example, I mean we're used to thinking about yield curves where we can put it in a two-dimensional graph. If I can use Black-Scholes I can actually use this implied volatility surface and it's a surface, it's not just a two-dimensional thing. I look at the surface and I'm thinking I'm going to have to know what this is at every future point in time, how am I going to do this? There are various ways of doing this. I can just keep it constant. I'm sure people have some up with other methods, or I could model it stochastically.

I'm moving to dynamic strategy and that's a lot more difficult than the buy-and-hold strategy. As David pointed out, dynamic hedging is typically done by matching Greeks. That's where you get the nice charts with matching delta and rho and vega. People come up with interesting terms for these things. What they mean is they match slopes, basically slopes of assets, sensitivities and liability sensitivities. I need to simulate the hedge strategy in this future financially fixed scenario. I'm going to talk about two main ways of doing this, a direct way and an indirect way.

The direct way is I think the more intuitive one. I'm just going to simulate what I'm doing. I'm going to simulate my hedge strategy, so at every future point in time, I'm going to compute whatever I need to figure out what my hedge position is, which is typically the Greeks for the hedge assets that I have and the Greeks and the hedged liabilities. Does that mean each future time point, at the beginning of the period, I compute my asset-liability Greeks? I determine what hedge assets I need to hold. I go to the end of the period, I look at what the value of this hedge. At the end of the period, the difference gives me my income and I recompute my Greeks and I rebalance my hedge positions. That's in reality how these programs work.

The indirect approach is like the buy-and-hold case. I guess I'll use the work assumption laden if that's a word. What I'm doing is I'm computing the value of the option that I'm hedging for each accounting period, the value of the option I'm actually hedging of the liability that I'm hedging. What does that imply? That definitely defines my hedge strategy because I'm saying I'm replicating, I'm reproducing, I'm manufacturing the entire change in that option value, in that hedged option value, and it assumes that I do so perfectly which I'm sure no one ever does. Isn't it nice to hear we're sure of that?

I'm going to talk about just a few issues that are related to doing these Monte Carlo valuations with a risk-neutral model. The first one, if you use your risk-neutral model, you have to have a risk-neutral model. At each future point in time you have to make sure that your risk-neutral scenarios are consistent with your realistic scenarios. It's also important to note and a lot of times people forget that if I compute Greeks, I not only have to compute one option, I have to compute a lot of them. That's run time. There's always the question of how many scenarios should I be running?

What's also important is if I say I hedge equities and interest rates, my stochastic model has to model equities and interest rates stochastically. Building this risk-neutral scenario generator has lots of issues associated with it that are certainly outside the scope of this, and given the time definitely. Time step is a huge issue for dynamic replication. Dynamic replication arguments are essential to balancing. In practice we can't do it continuously, but typically people look at the hedge positions at least several times a day and projections for capital pricing typically get done on a monthly or even courser basis. What is key is that it's the path of

equities and interest rates or whatever you're modeling stochastic that matters, not just the beginning of the period and the end of the period. Just to drive that point home, consider the error of replicating a fail to hedging an option over one month with a daily time step. That's what delta hedging does and these are idealized conditions. If someone says time step doesn't matter or error of delta hedging doesn't matter, they're wrong.

Hubert already talked about run time, and I'll follow up on his comment about age. He thought that we're not going to be around to get results. The other big issue is how to figure out hedge effectiveness. What does hedge effectiveness mean? How will I alter my future hedge cash flows to achieve the objective of the strategy? There are lots of sources of error. The first one is what I call the true hedge error. Dynamic replication is based on continuous time math. Well we can't do continuous time math in reality, we do discrete math. There are lots of assumptions that are met that lead to errors, and errors that are common to all or I should say most. Maybe somebody found the perfect model. Time step, I talked about that a little bit. I always have the wrong model. No matter what you pick it's going to be wrong. I will always not get the assumptions right. Market gaps can be a big issue. No model or that's no model that's hedgable or that I can hedge with anyway that I know of can model market gaps, plus execution is never going to be optimal. I can estimate that error directly using math or I can indirectly looking at back tests and get a feel for and estimate what that error could possibly be.

There are other two big parts about hedge effectiveness. Ineffectiveness is one's basis risk. Policyholders invest in all kinds of funds and typically I can't hedge them directly. Hedge instruments are only available on popular indexes. Again I can estimate future mismatch using historical analysis. I think the actuaries are well positioned to do that sort of analysis. Also my hedge is only going to be as good as my assumptions about the liability that I put in, that includes mortality lapses, exchange behavior. All normal assumptions get leveraged into the hedge calculations.

I've talked enough about what we can't do, so what can we do if we have to do this Monte Carlo simulation, if there are no closed form solutions? What can we do to make things better in some sense? One is I can speed up the Monte Carlo, the risk-neutral valuation. For example, there are several ways of trying to do conversions. I can run more scenarios for some policy and fewer scenarios for other policies that are in some sense not important for the hedge. For example, if I have a policy that's invested all in fixed and the benefit is way out of the money, I probably don't have to run whole scenarios for that policy. I can use good random numbers. Good means, for example, low discrepancy sequences that get rid of some of the correlation problems with normal random number generators. I can also simplify my hedge model. I think the initial bias of most people is going to be whatever I use in practice I'm going to use the same hedge model to simulate hedge strategy, and then they tell me there are other mathematical ways of speeding up Monte Carlo.

The next one is don't use Monte Carlo at all and don't simulate inside the capital projection, which means that I have to figure out what the hedge costs and the benefits are independently outside of my projection and it kind of sounds like reinsurance. I pay something for it and I get something back. Then there are mathematical alternatives. Most of them revolve around solving stochastic differential equations, things like finite differences come to mind. To do that you probably need to have a Ph.D. in math from MIT.

What do I like about reinsurance? It's fast. It's easy to understand. Most projection systems know how to do it. I start with the current cost of the offset I'm hedging. I turn that into periodic charge. Then the reinsurance itself will offset whatever I'm hedging, so if I go back to my accumulation benefit, the 10-year benefit, the reinsurance would simply pay that cash flow. I can incorporate hedge effectiveness easily by covering only a certain percentage of the hedged liability cash flow. What you can do to get yourself comfortable with that sort of an approach is if you're in the fortunate position that you can actually proximate your strategy with a static hedge, you can run that using a closed form solution like Black-Scholes and compare results.

Of course there are tons of other issues that I haven't talked about. Building or buying this hedge simulation capability is probably a question that's not answered all that easily. Good computing I believe is definitely going to be necessary to do any of these calculations halfway decently.

MR. HOPEWELL: Everybody can leave the room and know what capital initiatives are afoot in the world, what stochastic on stochastic means to them, what the choices are in calculating their capital both for VAs and beyond, and there were a few other things that we hope we've passed on or at least raised the profile of so you won't be surprised when you see them in the future.

FROM THE FLOOR: One of the biggest sources of ineffectiveness is model error. How do you model model error?

MR. HOPEWELL: I'll take that on from a qualitative perspective. You know what the components of your model error are.

FROM THE FLOOR: I don't think you do because you're saying that part of your model is right and part of it's wrong. How do you know which part is right or wrong?

MR. HOPEWELL: To paraphrase one of our government officials: the known unknowns and the unknown unknowns and the known unknowns are subject to statistical estimation. That is you can draw a cone of uncertainty or a bell curve around them. That's one component and then if you're a regulator you say you have to multiply by 1.5 to get the answer for the unknown unknowns. That's the

approach that's been taken in the past.

MR. MUELLER: Another way might be to look at your past effectiveness of your actual strategy. You can back test it. The market has been stable enough the last five years that back testing different strategies can be quite valuable in evaluating which strategy you want to use before you actually start going live. You can then use that ineffectiveness.