RECORD, Volume 31, Number 1*

New Orleans Life Spring Meeting
May 22–24, 2005

Session 25 Seminar

Annuity Risk Management Seminar: Stochastic Modeling for Annuity Risk Management

Tracks: Product Development, Risk Management, Investment, Financial Reporting

Moderator: Noel Abkemeier

Panelists: Noel Abkemeier
J. Timothy Gaule

Summary: As modeling techniques have evolved, the use of stochastic methods has been applied to the challenges of annuity product pricing and risk management. In this session, the presenters cover various subjects, including techniques and interpretations of stochasticism in fixed-annuity pricing, adaptation to equity-indexed annuity pricing and hedging, pricing of variable annuities (VAs) and derivative-based benefits (GMDBs and VAGLBs), dynamic hedging of derivative benefits, C-3 Phase I RBC for fixed annuities, and application and significance of C-3 Phase II and VACARVM for variable annuities.

MR. NOEL ABKEMEIER: I'm with Milliman in Williamsburg, Virginia. I specialize in annuity product development. Tim Gaule is vice president and valuation actuary and appointed actuary at Security Benefit Life. His comments will be heavily focused on the financial reporting aspects, something for which he's well-qualified.

We are going to go through the life-cycle flow of stochastic techniques in annuity products. We'll start off with a little overview of why, and then we'll look at the beginning of the game, which is the pricing area. We'll move briefly into hedging, and then finally we'll move into the financial reporting that continues on after that.

Frank Sabatini mentioned 12 risks facing the annuity market. What I found interesting is that an awful lot of them are the result of management action or

*Copyright © 2005, Society of Actuaries

inaction, business decisions and the unknowns. What's going to happen if the interest rates spike? How will the consumers respond if there are prolonged low interest rates? What about other consumer actions and utilization within the variable annuity products? There are a lot of unknowns. As we look at our use of stochastic processes, it seems as though we're only in a limited fashion really solving the risks. We're using these items more as a tool to quantify the risks so that we know what we're facing. In the process there are unknowns, which are the assumptions that you're putting in, and you have to live with those in your modeling process. You do the best you can, but realize that you may not hit every nail on the head. Again, with stochastic methods we are dealing with quantification and profiling of risk, and there is some degree of protection against risk. Hedging does that in advance, and then continuing, you are setting up adequate capital to make sure that when the rainy day comes, you're reasonably well-covered.

I will start at the beginning of the road with stochastic pricing. I want to address declared-rate fixed annuities, then equity-indexed fixed annuities. I do emphasize that equity-indexed annuities (EIAs) are fixed annuities. You will see that a lot of what comes out of the modeling process is going to show that the profile of an equity-indexed product is much like that of a declared-rate product. Then I'll roll into variable annuities, both the base product and the various derivative-based benefits added to them.

Why do we want to use stochastic methods in fixed annuities? The first reason is that we want to understand the range of the risk. We want to know how good is good and how bad is bad, not just what is average. We also, in the process, want to recognize the power of the renewal crediting methodologies. In particular, if you're using market-responsive interest rate setting methodologies, you want to factor that into your evaluation to determine how much on average it improves your situation.

Within fixed annuities, the major risk is the market value loss on the investments at withdrawal. Tied in with that is the need to properly or reasonably reflect consumer behavior, because the consumer behavior is assumed on average to be anti-selective, and so you want to mate consumers' anti-selectivity with the market value risk that you're facing.

The use of stochastic methods on book-value products is the most informative use because you are capturing the market-value loss on investments. That's the key item. On the other hand, if we're looking at market-value adjusted products, we find that stochastic methodologies are less informative. They tend to tell you the degree to which your market-value sensitivity is reflective of what is happening with your underlying assets and, as a result, it tends to show that you have a fairly low dispersion of profitability results. It's not as informative as on your book-value products.

In putting the model together, the first step is that we're reflecting historic interest rates, which reflect the pricing environment. The choice is, do we want historic or

do we want risk-neutral? To the extent that we're reflecting the real world in the model, historic interest rates seem to fit the bill. The renewal crediting formula is very important. Again, is it a market-rate responsive formula or not? To what extent does the function operate?

The dynamic lapses are a key item. Now, as I mentioned, questions were raised in an earlier session as to what consumers will do if there is a significant interest rate spike or if the low interest rate environment continues. We don't know the answer, but we have to do the best job we can of modeling the approach. There are methodologies in the market. You've probably all seen the arctan formulas and some exponential formulas. Whatever formula you use, the main idea is to move the interest-sensitive lapses into the right neighborhood. Exact precision is not necessary. It's desirable, but if you don't have exact precision, you'll still do well with the formulas that exist.

Stochasticism is a tool for determining liability duration, which can help you in your choice of the assets that you're selecting, and good asset/liability management will help control the risk that you're facing.

Finally, and perhaps of greatest interest to me, is that you have more informative profit measures. Deterministic pricing gives you just one simple snapshot, but stochastic pricing gives you a good range of results. In particular, we're looking at, say, the internal rate of return (IRR), or ROI or ROE. In any of these, it's of great value to look at two dimensions. One of them is your aggregate return, and this is distinguished from a mean return. The mean of your IRRs tends to be an unreliable number because IRRs are quite volatile and you can have outliers, which could greatly distort a mean. Consequently, looking at the aggregate is a better way to go. The second dimension is looking at the probability of exceeding a benchmark. You may be pricing your product to a 12 percent IRR, but it's of interest to know the odds that you're going to get at least an 8 percent or at least a 6 percent IRR or something of that sort. You want to have some measures such that your management team has a feel for how things might go and for what situations they should be prepared.

Let's look at other measures such as the profit margin and return on assets (ROA). In those cases, the mean is a good measure because both of these measures are continuous in their distribution and, therefore, quite reliable. Also, looking at a distribution of results becomes quite important. You may say, in general, that you expect 40 basis points of after-tax statutory ROA. Well, it's good to know and it's good for management to know the probability that, in the worse case, you're going to have zero profitability, or that you're going to have half of what you're shooting for. This distribution, I think, is good for communications within your whole management team, so everybody is playing according to the same rules.
As we take one step aside to equity-indexed annuities, I want to re-emphasize that EIAs are fixed annuities and if they're properly hedged, they will show financial profiles quite similar to a declared-rate product. There is the qualification of "if

properly hedged." The simplest one is if you have a static hedge. You know what you paid for it. You know what it's going to pay off. It gives you predictability. That gives you the greatest correlation with a declared-rate product. If you use dynamic hedging, you do introduce another dimension of uncertainty, because with dynamic hedging, to some extent you don't know your hedge cost until the hedging period is over, and how you do it can determine how stable that cost is in the tail end.

Determining your hedge cost is a case where you have to bring in your arbitrage-free basis of interest stochasticism. That is the environment in the derivatives markets, and if that's what your hedges are going to cost you, you better reflect the techniques that are used in that situation. In your modeling, a static hedge can be valued externally or internally if it's Black-Scholes-determinable and recognized directly in your model. If dynamic hedging is used, it becomes necessary to find ways to quantify the cost of the hedge and then bring that into your modeling process. Otherwise, you are facing some stochastic within stochastic, which is sometimes a good runtime challenge for the modeler.

Next, in looking at the experience of the product itself, we should be looking at the index on a historic basis as opposed to an arbitrage-free basis, because that is the real world in which we're living. Of course, the difficulty here is, what is a historic basis? We live through unusual situations: the bubble market of the late 1990s and the crash after that, the movement over time into a more international economy, different degrees of competition can affect how stock markets move and even the government deficits can well have an impact on what's going to happen in the future. You will have to make your own choice of what are true historic results or profiles that we're trying to recognize. It happens that if you're not precise on your historic model of how the index moves, it's not as large a problem as with some other products, because it is simply driving the growth rate of the product that you're projecting (how high is the account value), which is not as severe a problem as you might face in the derivative benefits in variable annuities, where how you project directly determines the cost of the benefit.

Within your equity-indexed annuity modeling, you should be dynamically resetting benefit formulas. This is assuming that the product is an annual-ratchet product, which is the dominant product in the market. Within your stochastic modeling, you do want to recognize how the benefit formula will change from year to year, which is a function of changes in the London Interbank Offered Rate (LIBOR) and in the index volatility. The index volatility, for example, can be derived from within the model itself. Look at the experience of the index that you're modeling, and use that with statistical methods to determine the volatility. Feed that back in and determine your option costs for your hedging.

You also have to recognize index-sensitive lapses. This is an unknown area, but you'll have to make some kind of judgment. Are index-sensitive lapses, which would be parallel to interest-sensitive, based on cumulative index experience? Are they based on current index experience? Are they based on expectations of what's

coming? The truth is probably that the lapses are most dependent on what people expect to happen, but you can't model expectations per se, so I think the result may well be that you will have to develop some kind of function that determines index-sensitive lapses in relation to the most recent crediting experience—the last one or two-years—and how you think people may react to that.

Again, I mentioned dynamically resetting the volatility based on a scenario. Another benefit you can get out of this is that you can test the impact of varying hedge ratios. Do you hedge your entire risk for all of the participants? Do you reduce your hedging in anticipation of deaths and lapses? To what extent do you reduce them? You can use the stochastic techniques and testing at various levels of hedging to see what gives you the most favorable results, the most stable results and so forth.

The next product is base variable annuities. Sometimes companies will think of variable annuities as a low-risk product because you've passed the investment risk on to the customer. Companies have learned that that's not anywhere close to the truth, particularly as the stock market went through its bubble in the late 1990s and then the crash in the early 2000s. Products that were sold in 1999 may have found their account values cut in half or worse. Soon the insurance companies realized how much of a risk is there and how dependent they are on the account values and the spreads on the account values to amortize their acquisition costs. In that respect, you find that a variable annuity is a much riskier product for the insurer than a declared-rate product, where you can change your interest rate and control some of how you're amortizing your costs.

Again, the large account value variances are important to be recognized through your stochastic modeling and, again, in the product. The reason why all this is important is partly that there are very narrow profit margins. The products have historically required low levels of capital, and low levels of capital require small margins to feed them, and that leaves little margin for error. It's very important to get a measure of what risks you're facing and, again, have management understand how bad is bad and how good is good.

In the variable annuity modeling itself, as we're projecting account values, you want to use a log normal regime-switching method to best reflect how the Standard & Poor's (S&P) goes or whatever other index with which you're dealing. Deep down for the base variable annuity you *might* be able to get by without regime switching, but when you're getting to your derivative benefits, you invariably need regime switching. If you're going to do it for that, do it for your underlying variable annuity, too. Of course, in the process you're creating, you want the returns for your multiple accounts to have good correlation among the various indices and the various fixed returns and so forth.

Within the modeling, you want to recognize economic capital. I'll say more about that when we get into the derivative-based benefits because those are the largest drivers of the C-3 Phase II RBC and the reserves under the variable annuity

Commissioner's Annuity Reserve Valuation Method (VACARVM). You need the facility within your model to recognize this one way or another.

In the area of market-sensitive lapses, this is, again, an area of unknown. You're going to have to make your best judgment as to the situation, but the challenge is how well people react in different market situations. What we saw over the last 10 years or so is such an unusual situation that it may not be indicative of future actions. You cannot tell, but you have to do the best you can to give a fair recognition of what the consumer behavior may be.

The profit measures that we're dealing with in variable annuities are essentially the same as those I mentioned on declared-rate products. The one change within it is that the IRR may be less important on the variable annuity than on the declared-rate products because of the low capital utilization, at least if there's no derivative-based rider on it. Because of the low capital utilization, you can have great volatility in the results. Because of that, on variable annuities you want to focus on measures such as looking at the dispersion of your return on assets or your profit margin.

The next step is the derivative-based benefits. We have a broad array of put-based benefits and one call-based benefit. The put-based benefits include the guaranteed minimum death benefit (GMDB) and the various guaranteed living benefits, such as accumulation benefit (GMAB), withdrawal benefit (GMWB) and minimum income benefit (GMIB). All of these will be touched on in a moment. On the other side of the coin, under the call-based benefit, your death benefit is increased in relation to the growth in the contract value and is, perhaps, conceptually covering the taxes that would be incurred on the additional growth that you have.

As far as the characteristics of the derivative benefits, in your modeling you want to recognize them in aggregate with the underlying VA. That's particularly important insofar as some of the risks compound each other. If your account values are falling, that's bad news for the underlying VA. That's bad news for put-based benefits. You definitely want to pull them together and observe and understand the extent to which they do compound each other.

Modeling with fat tails is imperative on the derivative-based benefits. You do have to properly recognize the frequency of the severe occurrences. Again, the two-regime-switching approach creates fat tails that are felt to be a fair representation of the history of the S&P and other indices.

When we get to the hedging of these benefits, you have to use your arbitrage-free models because that is the reality of Wall Street and how the various tools for hedging are priced in the marketplace. The C-3 Phase II and its cousin VACARVM require stochastic-within-stochastic concepts. This raises modeling issues of how you handle this and how you get your runtime under control. We'll touch on that in a moment.

Finally, the presence of these various riders on the underlying variable annuity does influence the persistency of the entire product, which is thereby also influencing the persistency of your underlying variable annuity. The presence of these benefits improve persistency, which will improve the profitability of the underlying variable annuity, which helps cover some of the costs of the rider that you have added to the product. This is an area of great unknown, though. These various riders have a short history. The GMDB has a long history, but it's the one that has the least influence on persistency. The living benefits are, perhaps, nine years old, starting with the first GMIBs. We've seen a little experience, but we're still in the unknown as to how it's going to operate over the long run and under various economic circumstances.

Again, using stochastic within stochastic concepts take time. If you have a complex model with multiple cells, it could run forever if you try to do everything precisely. We feel that, at least for the time being, reliable results are possible by using a methodology of determining the size of your C-3 Phase II RBC or the size of your VACARVM reserve through an external model, say an Excel spreadsheet development of a grid. Look at grid points that are a function of how long to the maturity of the benefit and how much is the benefit in or out of the money at a given point. Develop ratios of the C-3 Phase II RBC to the account value at each of those points, such that when you do run your stochastic modeling, you can call upon the grid, interpolate appropriately, and reasonably accurately bring your capital requirements in without tying up your computers forever.

Also within the modeling, you want to recognize the cost efficiency and the effectiveness of your dynamic hedging. Dynamic hedging isn't perfect. It has various costs. Part of it is the administrative cost of it, but you also don't really know the cost of your hedge until all is said and done. Part of that depends on how interest rates and volatility move, and part of it is how responsive your actual methodology is to changes in the market. Do you act soon enough to limit the losses that create the hedge cost? You really find it is necessary to, again, model your hedging technique to realize what share of the hedgeable risk you really capture with the technique that you're using.

As we look into some of the specific derivative-based benefits, the first one is GMDB, the grandfather of them. For the insurer, you do get a good picture of your risk profile as with all the other stochastic modeling. The good news about the GMDB modeling is that you're using relatively predictable utilization, which is the mortality rate. Compared with other derivative benefits, I would say that people die on schedule, whereas on other benefits it's their choice, and it requires good judgment and maybe some guesswork to figure out the best assumptions.

The extent that the GMDB is in the money could encourage people to stick with their policy. I think it's a very weak motivator, since people don't like to think about death as a financial investment tool and so forth. Yet, you do want to recognize the sensitivity of the in-the-moneyness on persistency, but it is probably not the most

sensitive measure that we're going to be talking about.

The GMAB has the convenience that there's absolutely no uncertainty about the probability of utilization. If your customer has stuck around to the end of the guarantee period, it will be exercised for him or her, so that simplifies your modeling situation. The persistency improvement as a function of in-the-moneyness and years to maturity has to be designed. There is no real track record here. You have to, again, use your best judgment, perhaps how *you* would act as a prudent investor if you had a GMAB that was a certain degree in the money, where you know it's going to pay off in a year or two if the market stays where it is and there's a very high probability that you're going to stick around.

On the GMWB, which is, perhaps, the most popularly sold product right now, utilization assumptions, again, have to be made. That's a very difficult situation. You first have to look at the product and how it is being marketed. Is it being sold as an immediate annuity, such that somebody who buys the product is invariably taking out, say, 7 percent a year and is on that path? Or is it a rainy day protection, such that it isn't utilized until the account value has dropped and it's clear that the benefit will be in the money and then you start making your withdrawals, which take a long time to work their way through? You do have to characterize the product as to how it is being sold and how it is being used before you can figure out what assumptions are appropriate. Just as with the other products, the extent to which it's in the money will improve your persistency or not, and you have to make your best judgments there. Again, it's an area of unknown because of the relative youth of the product.

On GMIBs, another dimension gets brought into the picture. You have to include a stochastic generation of the interest environment, since that is one of the drivers of the cost at the maturity date of the benefit. The interest that you're generating should be correlated with the market movements that you're recognizing for the account values. Utilization, again, is a great unknown. Again, some of these products have matured over the last couple of years, so maybe we're learning something. Your best judgment on how you feel people would perform is important. As on the others, persistency improvement is a function of how much it's in the money and how much customers understand what they have.

Next, and finally, is the call-based benefit. The earnings-based death benefit is an interesting product from a stochastic viewpoint. Its call-based nature is an offset to the put-based benefits that you have in your other benefits. For that reason, it's a nice thing to have in your portfolio. Of course, the techniques for modeling it are essentially the same as for your put-based benefits. I view the effect on lapse sensitivity as fairly low. I don't think it's nearly as important as products that are protecting you from losing money.

Once we have measured what to expect in the product, the next step is doing something about it. Dynamic hedging is one of the things that can be done. It is an

important hedging technique, and it's something that you will want to cycle back into your pricing model to integrate the hedging with the underlying risk.

As we look at hedging, something to keep in the back of your mind is that companies seem to be performing their hedging for two reasons, which are not always the same. There are economic hedges, where you're putting a hedge in place to cover the risk that's inherent in the benefit. Companies do that with products that are accounted for under FAS 133, so that would be GMAB benefits and GMWB benefits. If you put in an economic hedge in place, it also is going to give you appropriate, realistic accounting, so full hedging of those is attractive to companies.

The GMDB is accounted for under Standard of Practice (SOP) 03-01, where the accounting technique does not reflect true market risk. There the hedging is often done for the purpose of moderating the financial reporting impact. In such a case you may find that you don't want to fully hedge the product. It puts you in an inconsistent situation: on the one hand, you'd like to hedge it in order to control the risk; on the other hand, if you fully hedge it, you may be introducing volatility into your financial reporting. You find yourself looking for a happy medium.

Dynamic hedging for an EIA can provide additional profit, because it's on average less expensive than static hedging, or it could improve the benefit to the customer. It's your choice as to where you want to spend it.

Dynamic hedging is common for VA derivative-based benefits. It's just emerging for equity-indexed annuities. My observation is that companies that are in both the variable annuity and the EIA markets are the fastest ones to go into dynamic hedging for EIAs. They would be the ones that try to reap the benefits of the process.

The process of dynamic hedging is to put all of your risks in one bucket, evaluate them and then take appropriate actions. In such a case you may be throwing put-based variable annuity-related benefits, call-based variable annuity-related benefits and EIA benefits, which are call-based, all into the same bucket and determining out of that what is driving your hedging process. But even though they're all in one bucket, you can allocate the costs to the source of the liability because they each are going to be carrying their own set of risk characteristics.
Some of the methodology of dynamic hedging is covered in a later session, but briefly, the methodology is stochastic modeling on a seriatim basis or grouping of homogenous business to determine the Greeks: your delta, the sensitivity to the change in the underlying index; the gamma, the change in the delta as a result of the change in our underlying index; vega (or kappa), the measure of the sensitivity to a change in volatility; rho, for interest; and theta, for time to maturity. The technique is to measure the Greeks on an arbitrage-free basis of modeling, because that is how the derivatives that you're using are measured. Do this on a daily basis to get a blueprint of where you are and a road map for how you should respond.

Then execute it on a daily basis. You'll find, again, that put-based and call-based benefits can provide a natural offset in the process, although maybe the offset is somewhat limited.

I have covered the process of how you can look at and use stochastic methods to measure your risk in advance, get an understanding of it, and take some action, perhaps through hedging. Now we want to move into how you then follow up in providing adequate capital to cover the situation that you're hedging.

MR. J. TIMOTHY GAULE: My company, Security Benefit, is located in Topeka, Kansas, and we're primarily a variable annuity company. Seventy-two percent of our statutory reserves are related to VAs. Of those, there's an older block, about 40 percent of our reserves, that's associated with a 403b product that's sold basically to teachers in grades K-12. It has a guaranteed death benefit. It's a six-year step-up that I'd categorize as being a pretty tame benefit. In fact, when we've done our C-3 testing, we've found that the capital requirement on that particular product is actually less than the reserves we're currently holding. That has helped us, and I'll get into that more when we discuss the C-3 modeling.

It's also a product that on the separate account side tends to be very profitable because most of the assets are invested in internal funds, where we have better revenue sharing. There are some risks with it. It has a dollar-for-dollar benefit, and also, about half of the policies have a high guaranteed interest rate of 4 percent or higher. We've been monitoring the business quite closely, and we've been worried the last few years that with those high guaranteed rates there would be an incentive to move from the separate account over to the general account, and we have not seen that. Even within the separate account, we don't see a lot of movement, but we do continue to get premium. These tend to be pretty small policies from their first issue, where maybe you're getting a $50 or $100 premium a month, because these are teachers and they probably don't have as much income as the average person. But I do feel that the block helps us a lot, for example, with the C-3 calculations.

Then we have about 20 percent of our reserves associated with newer products that are an unbundled type of design. The basic design there is that the death benefits return the premium. Then you have all kinds of options as riders you can add. These products tend to offer riskier GMDB benefits, such as roll-ups at different interest rates (3 percent, 4 percent or 5 percent), step-ups and annual step-up. There's an enhanced death benefit and an option to combine the benefits.

On the living benefit side, a few years ago we introduced a GMIB option. A couple of years ago we introduced a GMWB. There are also options to waive surrender charges. There are many ways that the policyholder or the broker could game the system against us. What's interesting is that for the products that are still in the 403b market, the typical policyholder is not adding any options there, so you've backed a pretty tame benefit.

We do have one version of the unbundled product that's sold by registered investment advisors (RIAs). There are more aggressive subaccount options available with this product. For example, a lot of the Rydex funds are available. Typically, these individuals have a model where they're trying to time the market and take advantage of the fact that the assets are inside the annuity, so they don't have to worry about the taxability of gains if they're making a sale today.

The challenge we're going to have in modeling this product is that, first of all, we want to make sure we map the funds to an appropriate index, as opposed to the S&P. But what we've seen so far is that they haven't always done the best job of timing. For example, the return on that product last year in the separate account was just a little over 2 percent, whereas the return on the 403b product, where they're tending not to move anything, was 10 percent, which is more consistent with the market. It's going to be a challenge for my group in modeling that because we may have to make an assumption about not only what funds they're in, but maybe the fact that they're making the wrong choices at times.

As Noel mentioned, I'm the valuation actuary for the company. I serve as the appointed actuary. My area is responsible for statutory reserve calculations and the GAAP reserves, primarily addressing deferred acquisition cost (DAC). A couple of years ago we also started calculating embedded value (EV). That process is a deterministic process, but we calculate the cost of the options stochastically. We run 1,000 scenarios, come up with the average cost of each of those benefits over those scenarios and then put that back into the deterministic model. It's similar to what we're doing with SOP 03-01 because the benefit reserves are calculated stochastically.

My area also does cash flow testing. We've done some preliminary work already on C-3 Phase II RBC. You are all aware that there's an exposure draft, and we're assuming that the effective date for C-3 is going to be this year-end. I mentioned, too, that we do stochastic modeling for pricing. What I'm talking about there is that in the last few years we have done some acquisitions of blocks of business. My area has been involved in pricing out those acquisitions. Again, it's more or less a deterministic model. But when we get to valuing the cost of the benefits, again, we do that stochastically, much like we do with EV.

We're making the assumption that VACARVM probably won't be effective this year-end. How many of you are going to be impacted by either VACARVM or C-3 Phase II? How many of you are expecting C-3 to be effective this year-end? It looks like about half of you.

I'm going to be focusing on the financial reporting use of stochastic modeling, but I'll probably have a lot of emphasis on C-3 since that's the hot topic right now. I don't plan to go into a lot of detail on C-3 because I know there are some later sessions regarding C-3. I think the key with financial reporting models is that you

want to make sure you're modeling assets and liabilities under stress. It's important to capture what happens to the reserve and capital in those stress situations. It's important to have a sufficient number of model points that accurately assess all major drivers of risk. For example, you want to make sure you're capturing the revenues. By revenues, I mean mortality and expense (M&E) fees or our revenue sharing. It's going to be important to be familiar with the agreements that establish the revenue sharing, and based on those agreements, make an assumption about the appropriate revenue sharing to assume in your model.

As I mentioned earlier, we do have internal funds and we've had those funds for a long time. We can look at the accounting and see the revenue we get from those, but I'm not sure that we have revenue agreements in place. If we don't, I want to make sure that we do get that documented because I would need that when we do C-3.

You obviously want to model disbursements and expenses. As we've already mentioned, disbursements can be dynamic, depending on what has happened under the various economic scenarios. Your model should model asset behavior. If you're talking about the general account, you want to make sure you're capturing the risk of default, prepayments and calls. On the separate account, it's not going to be possible to model all the various subaccount options you might have, so it's going to be important to group those and map them to an appropriate proxy. Again, I go back to the example I gave of the product that's sold by our RIAs. We want to make sure that we don't map those funds to an S&P index because that would not be appropriate.

It's going to be important to have your model reflect policyholder behavior. Again, there's some guidance in the risk-based capital (RBC) report. One thing to remember is that policyholders don't always react rationally. You might assume that if it's a down market, they're going to take advantage of a GMIB or a GMWB. But I don't think it would be appropriate to assume, for example, that no one lapses, because although lapses should be dynamic and related to what has happened to the market, there are going to be times where policyholders need the funds for other reasons.

It's going to be important to balance accuracy and runtimes. Ideally, the more model points you have, the better. In our case, we have over 230,000 policies that we're trying to model for EV, cash flow testing and SOP. We've gotten those model points down to about 9,000 model points. I have a good group of people working with me, and they've come up with some creative ways of reducing those model points. For mortality, we use a blended rate based on the distribution of males and females in the cell. I want to emphasize that we're getting down to the cell level. For example, we have a block of business grouping all the GMIB. We'll look at the distribution by account value of males and females, and that's how we'll come up with the appropriate rate for that cell. We're not really doing it at an aggregate level; we're getting down to the cell level.

I mentioned the tax-sheltered annuity (TSA) product that's a big chunk of our business. There are different guaranteed interest rates that apply to that block. When we were doing cash flow testing, the guaranteed interest rate was one of the criteria for setting up a cell. We decided to do some blending there just like we did on the mortality rate. We took the higher guarantees, combined them and came up with a blended guaranteed rate for that cell. We did the same thing with the lower guarantees. I think the way that we split the guarantees was that 3 percent and lower we called a low guarantee. Above 3 percent would be a high guarantee. So within the cells we actually came up with a blended guaranteed rate.

For both C-3 Phase II and VACARVM, the assumptions should be based on the actuary's prudent best estimate. Best estimate is the actuary's most reasonable estimate based on the available information. Prudent best estimate would include a margin for error. The margin should be smaller if abundant and reliable data is available. Obviously, it should be larger if there's less data available. If you're talking about capital, the margin should serve to increase the amount of capital you calculate.

Again, I might point out that there's some additional guidance in the March report. There are some methodology notes at the end of that report. There's one that's labeled "C3-03" that deals with policyholder behavior. Another one that's important is C3-04. It deals with setting prudent best-estimate mortality. That's quite a bit different from what was in the September report. I would encourage you all to take a look at that.

If you are using an SOP 03-01 model for GAAP, it should be based on realistic assumptions. It's important in setting assumptions to consider dynamic behavior. We've already talked about that. You would expect utilization and lapse rates to vary based on economic trends. As Noel mentioned, one of the issues is that most of the living benefits are relatively new and so there's not an abundance of information available. It's going to be important to do some sensitivity testing to get a feel for the appropriate assumption.

The C-3 report talks about the prudent best-estimate assumptions. With C-3 Phase II you're calculating the conditional tail expectation (CTE) number. You should look at those worst 10 percent scenarios when you are addressing the concept of evaluation of scenarios. You should be looking at what is the appropriate assumption in those scenarios.

The fact that these are new benefits does point out the need for additional industry and company experience studies. Obviously, the industry studies may be a ways off, but it's important to try to collect good company information. One of the things that are helpful to us is monthly valuations. We're actually running our reserves monthly on a seriatim basis. That seriatim file is very important to us as far as tracking movements within the separate accounts. It's another source of getting some information on how many people have a 5 percent roll-up on the unbundled

product or how many took a combination of a 5 percent roll-up and some type of living benefit. That's a good source of information but it's not the only source, because it's not going to capture what's happening on surrenders. It's important to have good studies that will give you information on your duration-specific lapse rates, both full surrenders and partial surrenders. On the partial surrenders we found it important to be able to determine what are surrender charge free and non-free, because that's important for your models to get a good fit.

I tried to put together a laundry list of assumptions that you should consider when you're modeling, whether it's C-3 or SOP or, in our case, embedded value. Obviously, persistency is one. An important point to mention about persistency is that policyholder dynamics can vary by the type of benefit. If you have a dollar-for-dollar benefit versus a pro-rata guaranteed minimum death benefit, people might, because of the availability of that pro-rata benefit, be inclined to take a partial withdrawal to take advantage of that benefit if the product is in the money.

We actually tracked that. A few years ago there was an article in *The Wall Street Journal* that made everyone aware theoretically that there's this dollar-for-dollar benefit. We did see a pick-up in what we might call our dollar-for-dollar surrenders, but it wasn't huge. What was interesting was that a lot of those people surrendered the rest of what was in the contract a few months later. You don't really know what's driving that. Maybe they needed a new refrigerator and then later they wanted to remodel the whole kitchen so they had to take all the money out. You just never know.

As far as the mortality assumption, the key there for VAs is that deaths could be under-reported. Often what happens on a VA when the person dies and they're not in the money, the spouse may take over the contract. I know the way that's handled in our administrative system. Let's say it was the husband who passed away. If I didn't know that happened, I might assume that the wife had that contract ever since it was issued 10 years ago. You need to make sure that's being captured. Maybe one way to get comfortable with your mortality assumption or mortality experience is to look at your experience for deaths that are in the money because, obviously, that should be captured because there was a death benefit that had to be paid.

As I mentioned, there's additional guidance in the March report on how to set up prudent best-estimate mortality. Basically, you first develop expected mortality curves based on the available data. Then you want to apply margins to reflect any data uncertainty. Finally, the expected mortality curves are then adjusted based on the credibility of experience. Again, the margin should be applied in the direction it's going to increase. If you're talking about capital, for example, it's going to increase the required capital.

The next item I have on my assumption laundry list is asset-based fees. The fees are related to optional benefits. The challenge is to balance accuracy with increased

model points. Again, I'll go back to the two products I gave an example of earlier. Our TSA product has the six-year step-up. There's not a separate fee for that death benefit. It's buried in the M&E charge, so that's easy to model. But when you get into the unbundled products with all these various rider options with different fees, if you start modeling some of those together, you've got to be careful that you're getting the appropriate fee. For example, if I have a roll-up benefit and options that are at, say, 5, 4 and 3 percent, but most of the business is at 5 percent, I'll probably set up a model point for 5 percent, but I want to make sure that I've adjusted the fee I'm reporting in my model, because typically the fee for a 5 percent roll-up is going to be a little bit higher than that for 3 percent. I'm not bringing in more fees than I really have.

My next item on my assumption laundry list is mutual fund revenue sharing. As I mentioned, there's additional guidance in the latest report. You can include the revenue sharing in your C-3 model if it's received and controlled by the company, there are signed agreements in place and it's not already accounted for elsewhere. This is probably good in our embedded-value model. I know we bring in revenue sharing. When we set up embedded value, we did take a close look at the agreements, but I think with C-3 we're going to go back through it because I want to be in a position where I can back up what we're doing. The point I'm trying to make is that I think there will be a benefit in doing that for other processes.

Another item on my laundry list is expenses. If you're talking about C-3 or embedded value, you want to make sure that you include all your expenses: administration, overhead and marketing. A benefit we had of the EV project is that it forced us finally to feel a little more comfortable about our expense units. The reason I say that is we also have a family of mutual funds that's reported as a separate company; it's not part of the life company. Now that we've done an embedded value project (we're calculating embedded value for everything—the funds and the insurance), it has allowed us to get a better handle on our expenses. It's our finance area that puts those expense numbers together. They tell me that when they go off to different areas, whether it's marketing or actuarial or IT, everyone knows what expenses aren't theirs but no one seems to know what expenses actually belong to them. With this process we did finally get that down, where at least we feel more comfortable about the numbers.

Starting assets is next on the laundry list of assumptions. This is probably more of a comment for C-3, but it's also true for cash flow testing. You want your starting assets to be equal to your statutory reserves. The C-3 report specifies that in projecting your model, your working reserve is your cash surrender value.

As I've already touched on, benefit utilization is another important assumption. How will it vary by market performance and the degree of in-the-moneyness? The assumption should probably be that the policyholder will take advantage of benefits in the money. You would expect this, but, as I mentioned before, there may be reasons that they still need the funds. Let's say the market goes in that downturn.

That person could also be out of a job and maybe because of that, that person needs all the funds and, consequently, lapses the contract. It's probably not appropriate to assume that everyone is going to utilize these benefits. You obviously should assume that you're going to see higher utilization of guaranteed living benefits. Also assume that you're going to see your lapses fall, but probably not to the point that you don't have lapses.

Another important consideration is the transfers between the fixed account and the separate account. It should be dynamic. In other words, you would expect more transfers when market returns fluctuate. For example, think of the 403b product I mentioned earlier, where you got a 4 percent guarantee. Say you got a couple of years of down market. To me, there's an incentive to the policyholder to move to the general account. We tracked that, as I mentioned earlier, and I can't explain why they didn't move for the last couple of years; it could simply be that the business is asleep. But I don't know that it's going to be appropriate for me in future scenarios to assume that doesn't happen, because that block of policyholders are in their late 50s, and it seems to me that as they get older there's going to be more incentive to move the money. One reason you want to be comfortable about your assumption for transfers between the fixed and the separate account is that it's also a way for the policyholder to lock in a benefit. In other words, if there has been a significant downturn and they move over to that general account, if the market recovers, you're not going to get the benefit of that because it's not in the separate account.

The model should also capture the fixed-account flows. I'm not only talking about transfers there; I'm also talking about premium flows. Obviously, that varies by what market you're in. If it's an IRA, there are probably no future premium flows. But if it's a TSA market, there probably are flows, and you want to make sure that your model captures that.

Fund mapping is obviously a key assumption that's needed to reduce asset classes in the model. Obviously, you can't model all your various subaccounts. It's going to be important to review the prospectus and the historical performance of each fund. That could be complicated if you have a limited history—we have a number of funds where we do have limited history—or if there are manager changes. I've been at the company for 10 years, and for one particular internal fund, I think there have been 10 different fund managers. It's going to be hard to figure out the appropriate assumption there.

With C-3 you're allowed to map to proxy funds. It's allowable for the proxy funds to be a combination of the indexes, for example, that are provided with the prepackaged scenarios. It's going to be inappropriate to assume that the fund will outperform the market over an extended period. When I say that, I think of that product that's managed by our RIAs, where there are more aggressive funds. I know that if I go out and talk to those people, they're going to talk about how they're going to be able to beat the market for an extended period of time, but I

just don't see how that's going to happen. I think it's going to be important on the subaccount scenario. It's probably okay to map them to a more aggressive index, but that index should also be a lot more volatile. So if it's a down market, everybody gets hammered a lot worse than with an S&P index.

The next item on my laundry list is the number of model points. Again, it's important to balance accuracy with runtime. It's important to have credible assumptions at the cell level. It's also important, with C-3 Phase II and VACARVM, to accurately assess the tail risk. One thing we look at is that we get the net amount at risk of the model points and we make sure it matches the net amount at risk in our seriatim valuation file. Our model-point files all start with that seriatim valuation file. We feel comfortable using that file. Obviously, that's reviewed every month since we do monthly reporting. The auditors validate it. For those of you familiar with Sarbanes-Oxley, there are all kinds of other requirements around that material. We also want to make sure we have enough groupings of in-the-moneyness to approximate the seriatim run. When we put the cells together, we want to make sure we're still capturing the proper degree of in-the-moneyness.

It's important to remember that a large number of model points requires more computer capacity. The software that we use has the capability of distributed processing. We've taken advantage of that. We have a block of computers used only for the processing. We refer to them as our "farm" computers. That's what IT called them. When we originally set up that farm, we were using older computers. We found out that that didn't work. This last year, we upgraded all those computers. They're brand new machines, and they all have two gigabyte of RAM on them. They definitely have a lot of capacity, and that has been very helpful to us.

It's important that the projection horizon be long enough to model all material risks. In our case, we started adding GMIB benefits about three years ago. Well, the benefit is not even exercised until the 10th anniversary, so we obviously need to run more than seven years to get a true feel for the impact of that. I think most of our models will go out at least 20 years. We project our GAAP models on the SOP until only 5 percent of the policies currently in force are still left, so it varies a little bit, but it's roughly 20 years. One problem with going too long is the credibility of your assumptions. It's hard enough to predict what's going to happen in the second half of this year, let alone what's going to happen 20 years from now. It's also important to remember that longer periods slow your runtimes.

Another assumption item is future premium flows. Obviously, those are dependent on fund performance, value of options, economic conditions and the tax qualification code.

There are some options for the interest rate for discounting and general account earnings with C-3 Phase II and VACARVM. If you don't use an integrated model, you can use the implied forward rates from the swap curve. If you are using an integrated model, with the prepackaged scenarios you have indexes available you

can use. You want to make sure you're consistent from year to year.

Also, the rates used to discount the additional asset requirement (AAR) need to be reduced for federal income tax (FIT). If you don't know what that means, that's the additional asset requirement that you'll calculate as you're running each of your scenarios for C-3. You add that to your starting asset to get the total asset requirement (TAR). When you discount that back, the discount rates need to be reduced for FIT. It's important, too, that any interest assumptions on fixed assets reflect defaults.

In your model you can factor in hedges currently held and, if you have a clearly defined strategy, the model can take into account the costs and benefits of hedge positions expected to be held. That strategy needs to be in place for three months. Your hedging benefit should be adjusted to reflect imperfections in the hedge strategy. Again, there's guidance in the report. The C-3 report, Appendix 10, has quite a bit of guidance on it.

With both C-3 Phase II and VACARVM the Academy has provided prepackaged scenarios and a picking tool, so you can use those scenarios or you can use your own. If you use your own, it's important that they satisfy the calibration standards that have been set by the Academy. As a general rule, you probably want to run as many scenarios as you can. If it's fewer than 1,000, the actuary is strongly advised to check sampling error.

I have a couple of comments on SOP 03-01. There are no prepackaged scenarios available. When you do run the scenarios, it's important that the mean growth and volatility assumptions be consistent with your DAC assumption. We use the mean reversion method, and we make sure that that mean growth assumption is consistent with our long-term rate. That's not the only option you have if you use the mean reversion method. There is a practice note for SOP 03-01 that addresses a couple of options. If you use the mean reversion method, another option is that the mean of your scenarios during that mean reversion period should tie into your mean reversion rate, and then after that you tie into the long-term rate.

We use 500 scenarios for our SOP calculations. We actually did run 1,750 scenarios and found the results didn't change that much, so that's why we stayed with 500. Each one of our products has its own DAC model, so we have quite a few DAC models. You can see that we have an incentive to keep those scenarios down if we can, which, again, gets back to balancing accuracy with runtimes.

I have a couple more comments about the prepackaged scenarios. The full set of scenarios satisfies required calibration. Keep the monthly rates in mind because you need to integrate it in your modeling software. If you're doing a quarterly projection, it may be looking for something other than monthly. There is a scenario picking tool that's provided, and it's recommended that you use it. You're not supposed to cherry-pick the scenarios. Whatever scenarios you're using, they must

be available in an electronic format for regulatory review.

Chart 1 shows the list of asset classes and points on the yield curve that are provided with the prepackaged scenarios. They've added quite a few more asset classes since September.

Chart 1

C-3 Phase II and VACARVM Prepackaged Scenarios

Developed for:

- 9 Separate Account asset classes, plus
- 10 points on yield curve

3-month U.S. Treasury yields	6-month U.S. Treasury yields
1-year U.S. Treasury yields	2-year U.S. Treasury yields
3-year U.S. Treasury yields	5-year U.S. Treasury yields
7-year U.S. Treasury yields	10-year U.S. Treasury yields
20-year U.S. Treasury yields	30-year U.S. Treasury yields
Diversified Fixed Income	Diversified Balanced
Diversified U.S. Equity	Money Market
Diversified International Equity	U.S. Intermediate Term Govt
Intermediate Risk Equity	U.S. Long Term Corporate Bonds
Aggressive or Specialized Equity	

With regard to aggregation, you can take the higher capital requirements of aggressive benefits and offset that with the requirements for a less aggressive benefit. As I mentioned earlier, we have this large block of 403b business where, based on our preliminary calculations of C-3, we actually don't have a capital impact. However, with some of the more aggressive products, we do. We can aggregate those, and that reduces the overall required capital. Aggregation allows for the effect of diversification by age, issue year, benefit type and fund choices. It is, obviously, more conservative to aggregate at a lower level.

With C-3 Phase II there's also an option to calculate the C-3 Phase I benefits for fixed. We're a company that can take advantage of that because with C-3 Phase I there's a stress test and a significance test, and, based on those tests, we've been exempt from doing the modeling. In fact, two years ago, when that first came out, we looked at the benefit of the modeling and we saw that it could give us a lower number, but from what I understood of the rules we couldn't do that. Now with C-3 Phase II we can go ahead and model the fixed component of our available annuities. Once we do that, we can't go back to the factor method.

I don't think I can overemphasize the importance of sensitivity testing, especially if experience data is limited. It's important to do that sensitivity testing because you need to understand where your risks are. It helps identify the drivers of risk. It helps validate the base model and the level of conservatism you built into the model. It does increase runtime, but it's necessary. Ideally, you should try to sensitivity test against all your scenarios. If you're running 1,000 scenarios, you should try a sensitivity test against those.

As far as the CTE calculation for capital, it's CTE 90, meaning that you have to look at the worst 10 percent of your scenarios. For the reserves, or VACARVM, it's CTE

65. Unlike Phase I, there's not an Academy tool to do that calculation. It's pretty easy to do externally in a spreadsheet, or it may be supported by whatever modeling software you use.

There are some practical considerations. As I talked about earlier, distributed processing where you can send those runs out among several machines is important for runtimes. Never use an earlier valuation date. For example, for cash flow testing many of us use September 30 to reduce year-end load. The C-3 document has a couple of suggestions for how you can reduce your year-end load by doing a lot of your work earlier and using something called the "interpolation method" or the "informed projection method" to come up with that model-point file you need for year-end to do the C-3 calculations.

Obviously, modeling software capabilities have a lot to do with the ability to model these products the way we want. A couple of years ago, we moved to a new software package. One of the advantages is its open code, so it allows us to get in. If we have a new benefit, I don't have to wait for the vendor to come up with the support for that benefit. We can get in and put the coding in to support it. Obviously there are a lot of challenges then, because there are a lot of moving parts when you have open code like that. We've tried to come up with some ways of controlling that.

If you're starting with your C-3 model, you may want to build it off an existing model. For example, we started with our cash flow testing model. Again, I'll point out the need for experience data, especially for newer benefit types, and it's important for your assumptions to be consistent.

On the topic of certification and documentation, it's going to be important to comply with actuarial guidelines and standards. Practice notes are available. One has already been published for SOP 03-01, and there's a draft out for C-3 Phase II. I would encourage you to take a look at those practice notes as soon as you can because they do give you a lot of guidance. Down the road, if there are material changes, you want to make sure you disclose those. Make sure when you document your assumptions that you have documentation in there explaining how you came up with those assumptions. Internally we've been working with our auditors to make sure we're in compliance with Sarbanes-Oxley. This will be a slam dunk once we get done with that because we basically have to document where everything came from. If we've got a default assumption, how did we get that? It's not just some number that the investment department gave us. What's the background or back-up for it?

MR. MARK D.J. EVANS: I have three comments and a question for Noel. I think regime switching is one of several ways to address the problems you were talking about. It's nothing really magical. There are stochastic volatility models and other techniques that will get you to about the same place.

On the C-3 Phase II, you mentioned that practice note. It's mentioned in the treatment of hedging appendix specifically, but you're not required to do stochastic on stochastic to get credit for hedging under C-3 Phase II.

You had mentioned that call-put offset. I'd suggest looking at your higher order of Greeks, if you want to try and do that. I think that a lot of times you're not going to line up that well.

MR. ABKEMEIER: I agree. It has limited effect. It has some, but it's not huge.

MR. EVANS: On the grid approach that you were talking about for doing stochastic on stochastic, how do you address your pay-offs to the customer that tends to be highly path-dependent under that approach?

MR. ABKEMEIER: The grid development itself is starting from a point where the customer is, and each grid point is based on following the 500 paths. It would seem to me that it is recognizing the path-dependent aspects of what you need for developing the C-3 Phase II. It's a function of remaining time the benefit has to run and degree to which it is in-the-money. Granted, it does not capture the specifics of the recent path. The only point of the history that it is capturing is the in-the-moneyness. That is some reflection of the path, but, admittedly, it's not a full reflection of the path that got you there.