2004 Valuation Actuary Symposium*

September 20–21, 2004

Session 40 PD
Considerations in Generating Economic Scenarios

Moderator: Douglas L. Robbins

Panelists: Ellen Cooper
David Weinsier

Summary: A fundamental step in stochastic asset/liability modeling (ALM) is choosing a generator to create an appropriate set of economic scenarios. In this session, panelists discuss theories and assumptions underlying scenario generators, evaluating appropriateness of generated scenarios, choosing suitable economic scenario generators and use of representative scenarios.

MR. DOUGLAS L. ROBBINS: I'm going to be the moderator for this session, but I am going to tee up the topic a little bit by talking about some very basic issues. I know that here in the room there's probably a range of experience with generating economic scenarios. There are people who are just interested in what this is about, and there are people who run their own generator and want to improve it. I'm here for the people who are just starting out, and then I'll pass the topic over to Ellen Cooper, who is going to talk in a more serious, technical way about how scenario generators work and the different things you'll be using them for as valuation actuaries. Then Dave Weinsier is going to go the opposite direction and say that if you run so many scenarios that it's going to dim the lights in your building, here's how you can attack that problem.

I'm going to start out with my introductory session that I've just called "Now What Does *That* Tell Me?" A student comes up to you and says he has done your scenario set. He says, "Here you go. It's what you asked for—a stochastic scenario set with a mean return of 10 percent." You think: What does that mean? Often you'll hear a quote like, "From 1970 through 2001, the Standard and Poor's (S&P) 500 Index (net of dividends) had an average return of almost 10 percent." You heard that

Note: The charts referred to in the text can be found at the end of the manuscript.

statistic and the student heard it, and he created a scenario set that has an average return of 10 percent. But what can we make of that? If you look at the actual annual returns, January 1 to January 1 for all the 31 years, the arithmetic average was close. It was 9.75 percent. However, the S&P 500's value at year-end 1970 was about 92. At year-end 2001, the value was 1,148. If I take the return on the index to the $1/31^{st}$ power, I get a geometric average annual return of only 8.5 percent. So what can I make of that?

This reminds me of a story. It's about right after the ark, when Noah was dealing with all the animals. They were multiplying and doing just fine like he told them, except for the snakes. He asked, "What's wrong, snakes?" They said, "Look, if you want us to multiply, chop down all the trees in those woods over there, pile them in stacks right here and make us a house." He didn't know why, but he piled the trees up and the snakes slithered in. A couple of months later he came back and there were all these little snakes. Noah said, "That's great, but can you tell me now what the point of that was?" The snakes said, "Well, we're adders; we need logs to multiply."

The issue is that one average here is an arithmetic average. It's like adding, only you're adding the returns together and then at the end dividing by 31. For the other one, you're multiplying one plus the return together and then taking it to a power. You get different results when you do that. Annual volatility creates that phenomenon. Sometimes you have good years; sometimes you have bad years. If you had an 8.5 percent return in every year, you'd have an 8.5 percent geometric return, but your arithmetic average would also be 8.5 percent. Clearly, if you have a 0 percent return in one year and then a 20 percent return in the next, that gets you to a total return of 20 percent. But two 10 percent returns in a row get you to 21 percent. The volatility between the zero and the 20 causes your overall geometric return to be lower.

That effect is more pronounced the higher the volatility, as is the rate of the change. Try it yourself with minus 10 percent and 30 percent. You'll see that it's even more pronounced. In fact, there's a decent approximation for the difference, and this is where you get to the idea of multiplying versus adding. It has to do with the lognormal distribution. It has to do with the way the mean of that works versus a simple mu. The difference between the mean and mu is sigma squared over 2. That's also often the difference between an arithmetic and a geometric average for a scenario set where sigma is the annualized volatility. Thus, the bigger your annualized volatility, the bigger the difference between the geometric average and the arithmetic average.

How does either type of average compare to a single level scenario? From a profitability perspective, the arithmetic average produces similar results to the level scenario with the same statistic. Think about it with an end-of-year load. Let's say you have a $10,000 premium and you're going to charge a 1 percent load at the end of the year. If there were no increase, you'd get $100. If there were a 20 percent increase, you'd get $120. If you average those, you get $110, and if you had a 10 percent return, you'd get $110. It turns out that the result just tends to compound over time. So from a profitability perspective, as far as your loads, you're closer to your level scenario if you get a scenario set with an arithmetic average that's the same. But from a net-amount-at-risk perspective, a level scenario is closer to the geometric average because it's a better measure of average fund growth over time.

You want to create your stochastic scenario set, so you do it with this correct sort of average based on your level scenario. No. What you really want to do when you create your stochastic scenario set is base it on a statistic that you've developed the same way. You want to be careful of that. If you develop your average based on geometric and it's 8.5 percent, you want your stochastic scenario set to have a geometric average of 8.5 percent. If you've calculated an arithmetic average, you want your stochastic scenario set to have an arithmetic average. You want to make sure that the statistic matches the historical pattern you're trying to replicate, if you're creating a real-world scenario set. That's what this is really about.

If you then want to use a level scenario proxy, it should just be based on the correct stochastic statistic. That would be the arithmetic if you're trying to replicate profitability, and if you're trying to get at a notion for how net amount at risk runs off, then that would be the geometric. It's not the other way around, where you first do a level scenario and then base a stochastic one on that.

The second thing your student tells you is, "Here you go. Here's a stochastic scenario set with an annual volatility of 20 percent." Generally, "volatility" in economic scenarios refers to the standard deviation of annual returns. That's the most common way I've seen it expressed, and it's certainly important for valuing benefits where a big sudden drop is the main driver of cost. It's less important for roll-up benefits or benefits with a substantial waiting period, because volatility tends to even out over time, even if it's high. Mean reversion can make that even more true—if successive returns are interrelated, such that a big drop implies in your mind that there's going to be a rise in the future to compensate, then volatility is even less important. For the latter types of benefits—waiting periods, roll-up benefits—some measure of long-term volatility may be needed as opposed to just annual volatility.

There is a third statement you might hear when you'll want to say, "What do I make of that?" These cash flow testing runs include a mix of stock and bond mean returns and volatilities. So now you've replicated a bunch of different funds with mean returns and volatilities—stocks and bonds, maybe some cash, international equities. I'm not going to get into diversification because it's not really scenario-related; it's a modeling question. I'll just point out that different things can happen when you mix funds. In the short run you can say the following: If your funds are uncorrelated, the total variance will be the weighted average of the variances. So if you're mixing half stock and half bond, the variances are 10 and 20 and they're uncorrelated, you're going to get 15. That's variances. You have to take the square root to get the volatility.

For funds that are 100 percent positively correlated, the total volatility will be the weighted average of the volatilities. That indicates that there's not a substantial savings for funds that are 100 percent positively correlated, and that illustrates the futility of mixing highly correlated funds for variance reduction. You really don't run into it. In other words, if you mix two different large cap equity funds, they might be both modeled identically and, except for minor differences, they're going to produce the same pattern of spikiness over time. You're not going to get a modeling benefit by modeling both of them. However, for funds that are 100 percent negatively correlated, the volatility will be the difference in the volatilities, and that's very useful.

Say you have those two funds I mentioned. They both have a mean return of 8 percent, but you're trying to reduce your volatility. If they're 100 percent positively correlated, you get a volatility of 15 percent. If they're independent or uncorrelated, you get a volatility of 11.2 percent, because you summed the variances, not the volatilities. You get a nice savings in volatility through the independence, but if they're 100 percent negatively correlated, the volatility of the mix will be only 5 percent, which would be great. You don't find funds like that very often. However, a case where you can create this is where you have a fund that's large cap equity and then you purchase hedges in large cap equity. Now you are creating a 100 percent correlation situation. We'll get into that soon with the other speakers. Short-term volatility will affect long-term considerations as long as there's not an intense amount of mean reversion.

Earlier I said that based on a mix of funds, you could calculate your weighted average volatility in the short run. When I said that, I didn't mean that the theory changes in the long run; it's just the weights that change. That's my last point. Unless you're assuming rebalancing by your policyholders, your fund weights will change over time, and that will affect the way the volatilities work out, no matter what your initial fund mix was. You can have substantial savings in volatility from including a mix of funds in your scenarios, but the value of diversification depends on the funds being relatively poorly correlated. That is my final point.

With that, I'm going to turn it over to Ellen. Ellen is a senior manager at Ernst & Young. She just moved from the Philadelphia to the New York office. She recently passed her Chartered Financial Analyst (CFA) Level III exam and is awaiting her CFA charter. She has 15 years of experience with her firm in risk and capital management, stochastic modeling, ALM and derivatives-based hedging.

MS. ELLEN COOPER: I am going to talk today about a couple of different topics. I'm going to start out fairly technical, so for those of you interested in basically a cookbook of how to generate a correlated interest rate and equity risk-neutral generator, we're going to quickly walk through the steps of that process. I'm then going to walk into something a little more practical. I'm going to assume that we've generated a scenario set. We're going to talk about what the Greeks are, and we're going to do a case study using a guaranteed minimum withdrawal benefit (GMWB) example of how the Greeks can look very different under different assumptions in your risk-neutral generator. Then I'm going to talk briefly about nested stochastics and show you an example of some alternatives under nested stochastic pathwise calculations.

I like to think of scenario generations in terms of four categories. They're really all hybrids off of either real world or risk-neutral, but for what do we use real world? This is a question that keeps coming up, despite the fact that we live in the real world.

This whole notion of risk-neutral is difficult to understand. The way that I try to explain it to myself and to clients and colleagues is that it's just the price. There's no discussion about it. It's the way that the options market prices, and if we want to be consistent with the options market, we need to use consistent methodology in order to determine a price. We use risk-neutral for a couple of different things. We use it, as I said, to figure out the value of an embedded guarantee in our liability product. That will allow us then to be consistent with the way that options are priced in the markets, which then goes to valuing liability Greeks. I think of them as under the same umbrella. Calculation of liability Greeks and valuation of embedded guarantees are really in the same place. We also use risk-neutral to calculate the market value of assets, especially if we have an asset with an embedded option in it. We use risk-neutral to calculate duration and convexity as well, which actually are Greeks. I'll talk about that later.

The real world is when we want to look at performance—what's actually going to happen in the real world. We'll do things like look at present value of distributable earnings (PVDE). We would be using a real-world scenario if we want to look at embedded value or if we're doing a pricing exercise.

Nested stochastics are an offshoot of both of these. I'm going to go into more detail later. Essentially what we do is start at a specific point in time, at time zero, and at time zero we need to know the price of something. We need a risk-neutral process to determine the price of the guarantee, and then we're going to move forward in the real world to the next time step. When we get to the next time step, we again want to know the price. If the markets went from an index level of 1,130 at time zero to an index level of 1,000 at time one, how can I determine the price of my guarantee? Well, you need a risk-neutral process at that point in time. That again would be considered a nested stochastic process, and the new generation of models—Moses, ALPHA—is able to handle this kind of processing.

We also use nested stochastics with what I call "realistic with realistic" nesting. That would be for things like C-3 Phase II, where you have a certain prescribed set of scenarios that you have to use that are essentially real world. They have regime-switching in them, but they're essentially real-world scenarios. The Standard of Practice (SOP) 03-1, Actuarial Guideline (AG) 39 or variable annuity (VA) Commissioners' Annuity Reserve Valuation Method (CARVM) should all be sort of in the same place.

I'm going to drill down a little bit. I'm going to focus on risk-neutral scenario generation first. Here are the input parameters with which we need to be concerned: a risk-neutral interest rate process and a risk-neutral equity process. We want to correlate the two where it's important, and we need some kind of process around random number generation.

When we talk about risk-neutral process, what do we mean? The first thing that's important to understand is that all asset classes are going to produce the risk-free rate as their return. Short rates are going to move up the forward rate curve. We've talked about the fact that the options market prices this way. It's also important to understand that the futures prices are replicated, but they're not necessarily consistent with expectations. We're going to lock in today, but tomorrow, expectations could change the way that the term structure looks and, therefore, change the way the forward rate curve looks. Also under risk-neutral process, we want "no arbitrage" condition. We also need to use this for pricing and valuation in order to calculate a fair value.

There are all kinds of risk-neutral interest rate processes. I'm sure many of you have them or have built them, for instance Heath-Jarrow-Morton (HJM). The essential form is a(t) dt + b(t) dz. The a(t) is the time-dependent drift; dt is the change in time. The b(t) is your volatility, and the dz is your random process. The bottom line is that this scenario format has to be parameterized and simultaneously reproduce current market prices.

What we typically do is use a HJM format that essentially follows this form for the fund return:

$S_i(t+\delta t) = S_i(t)\exp\{(u(t) - \sigma_i(t)^2 * 0.5)\, \delta t + \sigma_i(t)\, \varepsilon_i(t) t^{0.5}\}$, where

i indicates the fund,
u(t) = risk-neutral rate, in this case the forward rate between t and t+ δt,
$\sigma_i(t)$ = volatility of fund i between time t and δ t and
$\varepsilon_i(t)$ = error term from the multivariate normal ~ $N(0, 1, \rho_{i,i} \ldots)$.

In this case the volatility of the fund i is constant over time, so $\sigma_i(t) = \text{stdev}\{\ln(F_i(t+\delta t)/F_i(t)\}$ across all t.

So for all of your indices, and especially if you're going to use hedging, S(t) is intended to be an index, like S&P 500, Russell or NASDAQ. We're going to move through time by S(t+ δt), equal to the last period index, and in the exponential, your u(t) is your short-term risk-free rate as it's moving up the forward rate curve over time. In this particular formula, we've used a constant volatility measure, so sigma squared is constant pertaining to that index, and it's an implied volatility. I'll talk about that later. You have a dt there as well. For your error term, you have your volatility again, and your error term has a correlation component in it with the way the different funds correlate with each other. Basically, if you take this formula, if you have an interest rate process and you use the u(t) as coming out of your formulas and you have implied volatility, you essentially can generate risk-neutral equity returns. If you wanted to try to price a 10-year European put, you can use this process to determine that you can replicate an option price.

In order to calculate that $\varepsilon(t)$ in the second half of the equation, this is the cookbook of how you want to do this. Probably a lot of you already know this, but essentially you need to do Choleski factorization. We need to come up with historical correlations of these indices over time. There are many different ways to do this, but typically you want to look at 10, 15 or 20 years—some kind of long-term correlation. Correlation stays relatively stable amongst funds. Once you come up with your correlation matrix, you then need to ensure it's factorable, which means that you can multiply its inverse by itself to get the correlations. That's Choleski factorization. You want to factor the Choleski matrix, and then you want to multiply that matrix by random errors. When I show you the Greeks and shocks, you shock with the yield curves to calculate your Greek calculations. A trick is to make sure that you're using the same random error elements for each of your shock scenario sets, because it will cut down on the number of scenarios that you need in order to converge to your price. If you have random elements that are different for each of your sets of shock scenarios, you're going to wind up needing to do 100,000 scenarios or more to try to converge. We've tested it, and you don't want to do that.

Basically you multiply these two pieces together and this, if you go back to the equation, is your $\varepsilon(t)$. It's essentially the factorization of the correlations multiplied by your random errors. You have the matrix of that, and you multiply that by your constant volatility.

Just to give you some perspective, Chart 1 shows a correlation matrix based on, I believe, 15 years of history, looking at S&P 500, NASDAQ, Russell and Europe, Australasia and Far East (EAFE) indices. The factorization is right below. If you take the Choleski matrix factorization and you multiply the inverse of that by itself, you'll get the correlation matrix.

Let's talk about some of the other input assumptions. In this particular example we're using a constant implied volatility, so basically we made some assumption. Maybe we said we're going to use the implied volatility based on three-year price. Essentially what we would do is look at potentially like 100 percent strike, and we would determine the implied volatility based on prices as of the point in time that we're running the scenario generation. That would be the input assumption. We would use a flat implied volatility assumption here. We then take this to the term structure, but this is just the quick, simple way, the "I need something, my boss needs a risk-neutral generator by the end of the day so let me put this together really quickly and get something" way.

The better way to do this is to use a term structure of implied volatility (Chart 2). By going out to the options market and looking at exchange-traded options, you can back into the implied spots, but what the model requires are the forward volatilities. In order to get the forward volatilities, you can use the equation that is written above the table. You essentially square each of the spot volatilities for the periods and times that you have them, and you can solve for your forward volatilities. Your forward volatilities, the $\sigma_i(t)$, is now also time-dependent. Basically we've taken this constant volatility measure and we've now put it in a vector instead. In addition, when we no longer have an implied volatility that we can get out in the market, we grade to historical in the forward volatilities, which is what I think most of the industry is doing because of the fact that we have long-dated options. We take implied volatilities for as far as we have them, and then we grade to historicals.

The other thing that is important to note is the skew. Skew is essentially telling us that as the strike moves relative to the spot, our implied volatilities are going to increase from even what we said they were going to when we calculated the forward volatilities. We want to take that into account because we typically are concerned about tail scenarios, and this is saying that volatility increases in the tail. We want to account for that. I'm going to show you an example of with skew versus without skew and show you how adding skew can impact, even on the base fair value. What we're saying here is that at time zero when you're first generating your risk-neutral scenarios, you look at your level of strike. Your level of strike would be able to, for example, if the S&P was at 1,130, as your spot moves away from that, as they diverge, your implied volatility will increase.

There are a few ways for you to check your results. A good way to check your results is by looking at the geometric averages (Chart 3). You should be able to ensure that the mean, for example, of your one-year rate equals the average of your short-term implied forwards at the one-year point. That should be a calculation that's relatively easy to check. If you look at your S&P 500, the mean of that should be equal to your mu minus sigma squared over two. The error term on average is distributed normal zero one. The mean is going to be zero. If you go through each of these, you can say that on average you have reasonable results. You can just go through here quickly and look at mu minus sigma squared over two.

Also, you see that we've gotten the various shocks. This is getting ready to calculate to the Greeks, but you can check in here and make sure that you have all your various shocks right. In this particular case, shocking up 10 basis points, your averages are going up by 10 basis points. Interest down is also approximately 10 basis points. We're shocking volatility up by 1 percent across the entire forward volatility term structure, and so you can see what that's doing. When we twist volatility up, it's lowering the equity returns, which makes sense. Volatility down is the opposite effect. Again, shock up S&P, that's an instantaneous shock up of 1 percent in a first time period and then everything else remains the same after that. You see a shock up of 3.52 for the S&P 500 versus a 3.49 for baseline, and that's basically showing the effect on average of shocking the S&P up 1 percent instantaneously versus the baseline. It's a good way to check your numbers and make sure that you have everything moving in the right way, because if not, it will come back to bite you when you start looking at the Greeks.

Chart 4 is another look at the distribution. This is basically pointing out that, in addition to focusing on the mean, you also want to focus on your distribution. Again, the error terms are distributed normal zero one, so you should be able to look at my mu minus sigma, which is equal to 66.7, close to 70 percent and 30 percent, and mu minus two sigma is equal to close to 90 and 10. Check those numbers and ensure that you have distributions that make sense period by period.

Now we're going to talk about a case study. I am focused on VA today. In this particular case study, we're going to look at some liability-based cases. Here's an example of a GMWB (Chart 5). This particular GMWB has a benefit cost of 35 basis points. I actually used 50, so pretend that the benefit cost is 50 basis points. What we're guaranteeing is that over time the policyholder can elect to withdraw 7 percent of the benefit base per year, with a guarantee that over the lifetime they'll be able to withdraw a total of the return of premium. There's also a death benefit in here, mortality and expense (M&E), age bands and so forth, but the important pieces are that the benefit is 7 percent per year and it's a return of premium over the period of time.

This is a real-world example, so here we're using stochastic real-world scenarios. You can get a sense using real-world scenarios. In this particular example we're assuming a worst case. We're assuming that everybody elects 100 percent of the 7 percent per year, which is a severe scenario, but it will give you a sense of claims. Chart 6 shows that in the 1 percent we have a large amount of claims. You can look at the relationship of the present value (PV) of claims minus PV of premium as well.

Now we have our product. We've priced it in the real world and we're comfortable with it in the real world, but we have to figure out how much it is going to cost to hedge this thing and what the fair value is in the risk-neutral world. We're going to use some Greeks to look at the sensitivities of the liability to be able to understand how it relates and what we need to do to match it in the asset-based derivatives world. There are a couple of measures that we're going to use. This is not intended to be all-encompassing; there are many more Greeks. There are more moments in time that can be used, but these are the basics. What is delta? Delta is a value change due to the change in the underlying index. It's an instantaneous shock in the level of a particular index at a moment in time and then understanding how the value changes as a result of that. Gamma is the change in delta. It's the second moment, and it's equivalent to the concept of convexity. If you can picture in your mind the price-to-yield curve, it's similar to its curvature but it's in the measure of the equity and the instantaneous shock. Rho is what we all know to be duration. You take a parallel shock up and a parallel shock down of your interest rates, and you calculate what the fair value would be under those particular scenarios. You would use the exact same formulas that you do when you calculate effective duration. Convexity, you all know, is convexity. Vega is a level term. It's consistent with how we do rho in that you do any percentage, but a level percent shock in your term structure of implied volatility both up and down to calculate what your change in value is due to changes in volatility. Theta is the time value and how the value changes relative to time.

Chart 7 is an example. We looked at the GMWB example under a bunch of different scenario alternatives. The important thing is on the bottom, the time zero liability cost. Time zero liability cost is determined by looking at the baseline GMWB claims divided by premiums and multiplying that by the starting WB charge. I believe that this was actually done using 50 basis points, so I apologize for the confusion there.

In run one, we looked at no skew. We had term structure of volatility. We also assumed that the asset allocation was allowed to drift. We did not attempt to rebalance the asset allocation, which means that in the bad scenarios when equities are moving down, you're going to wind up with larger percentages of assets invested in stock as a result of the fact that your equity values are going down. That's a very important concept because when we have VA contracts, we tend to not think about whether or not the policyholder is going to automatically rebalance or whether or not they've chosen to automatically rebalance if there's that option, but it has a dramatic impact on the actual fair value of the guarantee. When we look at that, the time zero liability cost is eight basis points. Now we add the assumption that everybody is going to rebalance at the end of every period back to their initial asset allocation. Look what it does to liability cost. It increases it to 14 basis points.

Then we add skew. That's really going to impact, less so in the mean. You're going to see it more in the tail, but even at the mean, it increases the base liability cost from 14 basis points to 16 basis points. You can imagine what it does in the mean scenarios. For run four, we went back to that bad utilization assumption, assumed everybody is going to utilize and it ran over to the next line. But look what happens if everybody utilizes 100 percent of the benefit at all points. We wind up with 32 basis points of liability costs. There's a very wide range. It makes a product like this very hard to hedge, because there are so many moving pieces. I'm not going to go into all the complexities around the hedging piece, but it's important to be aware that some of it has to do with policyholder behavior, which we're not going to talk about today, and some of it has to do directly with the way that you've set up your risk-neutral scenario generation process.

Now we're going to move on to talking about pathwise risk-neutral generation. I think the best way to describe it is to look at my chicken feet diagram (Chart 8). We start at time zero. At time zero, we are calculating a stochastic set of risk-neutral scenarios. We discount them back, and we come up with the full set of Greeks. Then we move on, using experience assumptions, to node one, except for every stochastic scenario we're in a different place. At the end of each one of these scenarios, we go through again and do a stochastic set of risk-neutral scenarios to come up with the Greeks again, because we need to dynamically rebalance once again. We keep moving forward and forward. You can imagine the amount of runtime. For those of you that have VAs and are faced with C-3 Phase II and have hedging programs in place, this unfortunately is where we're headed. To solve the problems of runtime in this particular issue, we have to figure out how we can run our hedges through time to be able to get the appropriate capital relief for our C-3 Phase II, while not dimming the lights on our models and our computing power. That will be, I think, one of our big challenges in the next 12 to 18 months, or however long it is.

In the example I'm going to show you, what did we do? Again, we used realistic scenarios to project the experience, and at each point where we're going to rebalance, we go into our risk-neutral nested scenarios. There are a lot of moving parts. One of them is lapse rates, one of them is utilization and one of them is going to be changing capital market assumptions. We're going to recalculate the fair value at each point. In this particular case we're also going to recalculate all the Greeks, use those to dynamically rebalance to whatever algorithm we have on the asset side and move on to the next node.

Chart 9 is an example of some pathwise liability Greeks under alternative scenarios. In this particular example, on the top left is the liability value as it's moving through time. It's using the nested stochastic process that we talked about. These are two different deterministic experience scenario assumption sets. We're moving through a certain state of the world, and we stop and calculate a stochastic set of risk-neutral scenarios. Depending on the state of the world, you can have very different Greeks, very different rule sets and very different understandings of what we need to do in terms of being able to hedge these particular liability guarantees.

FROM THE FLOOR: How many different scenarios are you running at the intermediate projection points?

MS. COOPER: That's an excellent question. We're doing 1,000 at each of the intermediate points, as well as 1,000 shocks. We feel that we cannot get a good answer unless we're running at least 1,000 at each of those. We're playing around with a bunch of different techniques to cut down on runtime. That's a great segue into our next topic.

MR. DAVID WEINSIER: I'm with the Atlanta office of Tillinghast. I specialize in ALM issues, securitization of XXX and AXXX reserves, as well as embedded value and the life settlement market. Prior to joining Tillinghast, I was with ING for a number of years, focusing primarily on product management.

We're all doing stochastic modeling. In this day and age you have to. Whether it is a new regulation or a new guideline that comes out, whether you're pricing a VA, or maybe you have a new GMWB procedure and you're trying to determine the cost you're going to charge on that feature, you're going to have to use stochastic modeling to come up with that charge. Make sure it covers all your implicit costs at maybe the 95th percentile. Maybe you're using stochastic modeling to get a handle around your interest rate risk on your fixed annuity or universal life (UL) block. It has a lot of different uses.

Stochastic modeling, as we all know, is not new. It has been around a long time, but it has not been as virtually required as it is today. There are obviously a lot of upsides to stochastic modeling. It's almost essential if you're going to do a good job to determine risk, to measure risk, but it's not without its downsides as well, one of the biggest of which is the inevitable runtime issue. If you're running 250, 500, 1,000 or 5,000 scenarios, unless you have Big Blue sitting on your desk at work, you have a runtime issue. This is an issue that every company has to face. So what are you going to do about that runtime issue? You have a couple of options. You can ask your senior management to go buy you a big room of computers. That request normally does not come across very well, although I have seen that situation in some companies. But even if you have this big bank of computers, say at year-end you have this big corporate model you have to run and you have 1,000 policies, a reasonable block of business, and you probably have somewhere between 1,000 and 5,000 model points, that's going to take you a little time to run an ALM model.

What you do is you try to shrink it down, chop your model up to bring the number of model points down. Sometimes that works better than others. It depends on the type of product. For instance, fixed annuities have fewer moving parts than whole life. They're going to be maybe a little easier to consolidate. It depends on how homogeneous your block is. It depends on how robust your model needs to be. Cash-flow testing has different needs than maybe a short-term planning model. I've actually seen a company spend a lot of time in consolidating their liabilities. They did a great job of cutting down their model points 80 percent (or something like that) and then they went to run their model and discovered that it wasn't the liabilities causing the runtime issue, it was the assets. They were almost back at phase one.

Today I'm going to try to help you reduce your runtime via scenario reduction using representative scenario methodology. I think we'd all agree that there's pretty much a one-for-one relationship between the number of scenarios you can reduce and your runtime. If you reduce your scenarios from 1,000 to 100, your runtime is going to reduce about the same—90 percent.

Next I'd like to talk about the Tillinghast cash-flow testing survey that we did in

December 2001. We did get good participation on this one—61 companies replied out of 84. Among other things, we asked them if they were using a method to reduce the number of stochastic scenarios, and 26 percent reported actively using a method to reduce the number of scenarios, which we thought was impressive. Even more impressive were those that were using some type of methodology. They were seeing a 90 percent reduction in the number of scenarios. That's pretty good.

To reiterate, why is there the need to reduce scenarios? As we pointed out, running thousands of scenarios can certainly debilitate both time and resources. Additional, faster machines are a partial solution. Smaller, faster models may help a little or a lot. That gets back to my points on model consolidation, but reducing scenarios is going to improve your runtime—again, about a one-for-one—and your flexibility. When I say flexibility, say you have a big corporate model, an all-in model, that takes about a week to run or even a weekend. Say it has to be run over the weekend and you're at year-end; you're under a time crunch. It's hard to get this thing right on the first try. You're going to want to change assumptions. You're going to want to set up some new stress tests. If it takes a weekend to run, you're in trouble. That's going to really limit your flexibility. What usually ends up happening is that your time deadline hits before you finish up all the runs you need, and you end up basing your results or your decisions on models that may not be as fully flushed out as you would like.

There are some commonly used variance reduction techniques, in other words, methodologies to reduce your scenarios. You've probably heard of most of these, if not all. The first involves antithetic variables, in which for each random number x you're going to construct another scenario using 1 minus x. The scenarios are symmetric, which does preserve the expected return. One of the features of these antithetic variables is that they do a very good job of converging on the mean of the distribution, but they do lose the shape of the distribution. While they would do a very good job in calculating the price of an asset, they would probably not do such a good job when you're coming up with capital requirements, when you're really focused on those tails.

Importance sampling and stratified sampling also do a decent job. Low-discrepancy sequences are probably the most common, and in these it's the same story as the antithetic variables. They do a very good job of converging on that mean, hitting that expected return, but again, with C-3 Phase II and AG 39, you have to focus on those tails because you're calculating your conditional tail expectations (CTEs). These methodologies are probably not going to be optimal for your purpose.

What is representative scenario methodology? It's going to reduce the number of scenarios needed to apply stochastic asset-liability cash-flow models, with results remaining comparable to the full scenario run. That's the goal. We're going to give them some level of tolerance that you're going to have to set on your own—95 percent is not far off. You're going to want to take the number of scenarios from 1,000 to 100, maybe 50, and get the same results as you would by running the full set. This methodology was developed by a good friend of mine, Alastair Longley–Cook, and Yvonne Chueh, just a couple of years ago. This is not new, by the way. I didn't invent this. It was presented at the Canadian Institute of Actuaries' meeting. You might have seen a paper on this before, so this will be a refresher, but if not, it will introduce you to some slick methodology.

Let's talk about the inputs and outcomes of stochastic asset-liability cash-flow models. Scenarios are going to be your inputs. Those are going to be your stochastic inputs to your ALM model. This methodology is going to work for both interest rates and equities. That's important to keep in mind. It's not just for VAs with the guaranteed living benefits. It's also going to work on your UL or your fixed annuities in coming up with measuring your interest rate risk. Cash flows are going to be your outcomes. It's the creation of these cash flows that obviously is going to be time-consuming. In other words, coming up with the scenarios is easy. Anybody can be coming up with scenarios; it's the cash flows that take the time and the care.

Chart 10 shows a mathematical formula. You're going to see it again later, so let's not get into too much detail here. The point on this chart is that central to the representative scenario methodology is the notion of coming up with a distance between scenarios. You have scenario A; you have scenario B. You're coming up with a distance between the two. By coming up with this distance, this is how we're going to be able to get good representation in our tails. If A is on one end, and you're able to figure out that B is on the other end of your spectrum, you'll see in the results that that's going to keep you from these representative scenarios bunching in the middle, which is what you don't want. You want to get adequate representation out here in the tails. This distance formula and methodology is what's going to allow you to do that. As I said, representative scenarios are chosen so that they are far apart from each other.

Chart 11 points out what I just said. The scenario sampling algorithm uses distance (D) to ensure adequate tail representation. What we have on the left is 1,000 scenarios, cumulative distribution function for present value of future surplus, I believe. Going on your Y-axis is going to be cumulative probability. Your X-axis is your PV of future surplus. As you can see on the right, you have 50 scenarios. On the left is 1,000; on the right is 50. You can see we have good representation in the tails here, from both the top and, more importantly, on the bottom left. Again, if you were just using maybe a sampling technique, you would get a lot more bunching in the middle.

So why use representative scenario methodology? Again, you're going to reduce your runtime by reducing the number of scenarios, giving you the greater flexibility you need to run your model over and over. It's going to give you actuarial knowledge of the relationship between scenarios and cash flows. What I mean by this is that on each of your representative scenarios, you're going to have a probability assigned to it. It's going to be representing x number of scenarios in your big set. Let's use our example, 1,000 down to 100. It's not one representing 10. It's going to be one representing two, another one representing 20 and another representing 90. You're going to have probabilities on each of these representative scenarios that represent the entire set. This has been proven to be more effective than simply generating fewer scenarios or randomly picking them out. Results are often very similar to those using all scenarios, which, of course, is the goal.

I'm sure that you've heard a lot about regulations in the last couple of days. Why do we have to do stochastic? GAAP SOP 03-1 means you have to do stochastic. Risk-Based Capital (RBC) C-3 Phase II certainly does. That's going to be capital requirements calculating CTE 90. In fact, our case study coming up is going to be based on coming up with a CTE 90 for the RBC C-3 Phase II proposal. Canadian requirements, I believe, specifically recommend 1,000 scenarios. For AG 39, now we're talking reserves, VA, guaranteed living benefits, CTE 65 maybe, but again, stochastic is going to be required. Complex and path-dependent guarantees could require thousands of stochastic runs of the cash-flow model, and representative scenario methodology is going to help with this new runtime issue, which we're going to see later.

Here is the case study model. We're going to focus, as I mentioned, on RBC C-3 Phase II projected to be in place in December 2005. A VA model was used to calculate capital requirements. We did start the model one month after issue to simplify things, effectively ignoring the point-of-sale costs. Specifications and assumptions are based on typical VAs available in the marketplace. There is nothing fancy here. We're focusing on reducing the scenarios.

This is where it gets fun. Here are the five steps. First we're going to create the risk parameter scenarios. That means your big set—your 1,000, your 10,000, your 5,000. Then we're going to choose representative scenarios out of that big set. The third step is assigning probabilities to the representative scenarios. That's a new technique, which is interesting. Then we're going to, of course, run our asset-liability cash-flow model and lastly, analyze the results.

Regarding the creation of the risk parameter scenarios, we're going to have 1,000 equity scenarios based on the RBC C-3 Phase II version of the regime-switching lognormal model. These are going to be calibrated to the S&P 500, just like what's in the proposal. There will be 360 monthly returns (30 years). To keep things simple, we left interest and inflation deterministic, as well as credit spreads and default rates.

Chart 12 deals with choosing the representative scenarios. We're going to take the number of scenarios down from 1,000 to 50, a 95 percent reduction in scenarios, and we're looking for a 95 percent reduction in runtime. Here are your algorithms that Ms. Chueh suggested and on which we've done some testing. You have two labeled "D_1" and "D_2" and then one below labeled "S," called the "significance method." The first two are relative present value distance methods. These maximize the distance between your chosen scenarios and your other scenarios. In other words, r is going to be one scenario, scenario A. Then r' is going to be scenario B. V is the discount factor. In effect, all you're doing is calculating the distance between scenario A and scenario B by summing, multiplying and discounting in order to determine which ones are the farthest apart.

D_2 does the same thing as D_1 except that it does not have the discount factor V; it's implied in the formula. That's the one we're going to be using in our numerical case study coming up. The significance method is going to be a little different. It's going to be a little simpler. As you can see, there's no r', just r. There's just one scenario, A. What you're doing on the significance method is you're taking scenario A, duration 1 to 30, as you can see. That's going to be your k, and you're just coming up with a variable. If your first duration rate is 9 percent, you're going to take 1 over 1.09. That's going to be your first number. You're going to add that to 1/1.09 times duration 2 and so on and so forth. It's a pretty simple formula. It's easy to develop this on a spreadsheet; we've done it before.

On the significance method, after you run your 1,000, you're going to have 1,000 different numbers. All you're going to do is rank them one to 1,000 and pick out every 20th. The significance method is also known as the uniform method. It's almost like a stratified sampling where you're ranking them from top to bottom and you're getting good sampling throughout. As you'll see later, the significance method does a very good job of focusing in on the mean, but does not do as good a job on the tails as the distance methods.

The third step is assigning the probabilities. The probability p for a representative scenario—that's one of our 50—is going to be the number of scenarios assigned to it divided by itself. Another way to say it is that p is the probability of occurrence of any scenario assigned to the representative scenario. If your p on one representative scenario is 90 percent, that means there's a 90 percent probability of occurrence of one of your 1,000 being assigned to that scenario. The probability p is going to be based on the inputs, but applied to the outputs (cash flows). I'll show you how to do that. Now, this is what stresses the importance of a close relationship between your inputs (scenarios) and your outputs (cash flows).

We can break our steps down even further. Step one: choose an arbitrary scenario out of 1,000 and call it Pivot #1. Step two: you're going to calculate your distances, based on the formula we just saw, from Pivot #1 to the remaining 999 paths. This is easy to set up on a spreadsheet. Step three: name the scenario with the largest distance to Pivot #1. You've calculated all your distances, r and r', and the one with the biggest number you're going to assign Pivot #2. That's how you're getting the A and B really far apart. Step four: calculate the distances of the remaining 998 to both Pivot #1 and Pivot #2. Step five: assign each of the 998 to the closest of the two. Now you have two distinct, disjoint groups. Step six: rank these 998 again in descending order. You're going to rank the distances, and the one with the largest distance is called Pivot #3. Again, we're taking the one with the largest distance, making sure these scenarios are spread apart. Step seven: follow the above procedure to select the remaining 47 pivot points. Step eight: if the number of scenarios associated to a pivot scenario is x, then assign a probability of x over 1,000 to this pivot scenario. If you have 900 of your 1,000 assigned to Pivot #1, it's going to get a probability of 90 percent.

Say I have eight scenarios here, represented by eight black dots. Step one: choose an arbitrary scenario out of these eight. We're going to go from eight to three, by the way. That's going to be our case study. Choose an arbitrary scenario out of the eight and call it Pivot #1. Pick the one in the middle. Calculate the distance from Pivot #1 to the remaining seven paths. Name the scenario with the largest distance as Pivot #2. Now we have Pivot #1 and Pivot #2, nice and spread out. Step four: calculate the distance of the remaining six to Pivot #1 and Pivot #2. I'm picking the closest one. I'm going to assign three to Pivot #1 and three to Pivot #2. Step six: rank the rest in descending order. The scenario with the largest distance is called Pivot #3. Now you have to calculate your distances again. One is now closest to Pivot #3, two are closest to Pivot #1 and two are closest to Pivot #2. What do we have here? We have three scenarios—Pivot #1, Pivot #2 and Pivot #3. Now let's assign our probabilities. How many of the eight scenarios are assigned to Pivot #1? It looks like three out of eight, or 38 percent. How many of the eight are assigned to Pivot #2? That's also 38 percent. Pivot #3 has two, so that's 25 percent.

We've gone from eight scenarios to three with probabilities. Pivot #1 is going to use a probability of 38 percent, as is Pivot #2. Pivot #3 is going to use a probability of 25 percent. That means two scenarios are in Pivot #3.

The fourth step in the methodology is to run the cash-flow model. The cash-flow model was run for the representative scenarios, generating 50 outcomes. Each outcome has an associated probability, just like what we came up with. The outcome used for RBC C-3 Phase II is the worst PV of year-end statutory surplus. Metrics such as required capital can be calculated directly from the set of 50 outcomes and probabilities.

Let's take a look at some results. Are you all familiar with how CTE works? You take your PV of surplus at the end of each duration, and you take your biggest one for each scenario. Then you line up your 10,000 scenarios, pick your bottom 10 percent and you take the average. That's about right. Let's say we had 1,000 scenarios reduced to 100, and you have your probabilities. At the end of the day I still want 1,000 outcomes, but we've only run 100 scenarios. I want to take your probabilities and multiply by 100. For example, if one of your representative scenarios has a 90 percent probability, I want you to use that one 900 times. At the end of the day you're still going to have 1,000 sets of results. For one of your representative scenarios, if your probability is 90, you're going to have 900 of those when you go into calculating your CTE 90.

Chart 13 has two lines, one light and one dark. This is again a cumulative distribution function. On the X-axis, same as before, is worst PV of surplus as percent of account value, again with cumulative probability on the Y-axis. The light line is going to be 1,000 scenarios. Your dark line is going to be your representative 50. You see the spots there. As you can see, this does a very good job of giving you representation in the tails. Again, if you were just doing a stratified sample, you'd see those bunched up in the middle, but here that's not the case. Because of this distance method and our probabilities, we're getting these representative scenarios way out in the tails.

Chart 14 is the same picture, but focusing in on the tail. See how the top of the Y-axis is just 10 percent. Just focusing on the tail here, this is an excellent representation. This is X-axis, worst PV of surplus, so we're actually doing the CTE 90 calculation here.

Remember the significance method, the one that does a good job of getting you to the mean, but doesn't do such a good job on the tail (Chart 15). Our dark line on this one is going to be the same set of results using the significance method, not the distance method. It misses the tail. It only stretches as far as that negative half percent maybe. You're going to miss your 99th percentile, your 98th percentile—the areas that you have to make sure to catch when you're determining capital requirements or reserve requirements.

I do want to point out one thing. We recently had a chance to do some of this work. It was for a project involving compliance with GAAP SOP 03-1. While representation in the tails was important for the project, that GAAP statement of position does state explicitly that guaranteed minimum death benefit (GMDB) and guaranteed minimum income benefit (GMIB) costs need to be generated using scenarios that are consistent with the single scenario used to amortize deferred acquisition cost (DAC). In other words, it needs to reproduce the mean very well. We started with the distance method, just like we did today, but we couldn't get the distance method to reproduce the mean. So for this particular project, while the tail was important, it was not as important as representing the mean. We did end up using the significance method for this project, reducing the number of scenarios from 1,000 to 50, 95 percent, and the results turned out just fine. I don't mean to throw out the significance method; it does have purposes. GAAP SOP 03-1 is one of them; RBC C-3 Phase II may not be one of them.

Let's sum it up. This methodology appears to be both accurate and robust based on the theory behind it, as well as several opportunities that we've had to do some real-world work on it. Accurate CTE 90 RBC C-3 Phase II values were calculated with a 95 percent reduction in scenarios and runtime. Keep in mind that this could be used for many purposes. You don't need a new regulation or new law to use this methodology; all you need is a runtime issue. You can use it for economic capital determination, reserve setting, risk analysis management, ALM, pricing, optimizing investment strategies or hedging strategies.

MR. KONRAD P. SZATZSCHNEIDER: Our scenarios have not only an equity component, but also a correlated interest rate component. We use the representative methodology, but we had trouble deciding how to weigh the equity and interest rates when picking scenarios. Now you have two pieces to the distance formulas, so I was wondering if you have any opinion on how to weigh each.

MR. WEINSIER: That's a very good question, not only in terms of representative scenario methodology, but even outside of that. Often the question we get is: For a given interest rate scenario, what kind of equity scenario ties with that? You're saying you used the methodology and tried to tie the two together. That's not something with which I have worked. That's a very good question. What if you calculated them independently? Do you feel that would skew the results a lot?

MR. SZATZSCHNEIDER: What scenarios would you pick? What do you mean by "independently"? Let's say I have 1,000 scenarios, and each one has 30 years of equity returns and 30 years of interest rate returns. In our methodology, I would pick out 50 equity scenarios and the corresponding interest rate ones.

MR. ROBBINS: Let me give you an idea from which you can extrapolate. I have no idea what kind of benefit you're valuing, but one place I can see that coming up is if I'm doing a GMIB where you issue a VA and no matter what the policyholder invests in (international equity, super risky, whatever), you guarantee he can annuitize after seven or 10 years at some guaranteed rate. The equity component is that the fund might go down. The interest component is that I'm expecting a spread on my payout annuity to make up for some of what I lose, but I might not get that spread if interest rates have sunk way down. Maybe I do my distance using the equity rates but adjusting them for any loss I get if interest rates are also low. Now does that help you if you're doing something different? I don't know. Maybe you just have a separate account and a fixed account and are modeling both, but if there's some way you can come up with measuring the cost of the interest rates compared to the cost of low equity rates and assigning a proportion, I think that's what you're trying to do. At the end of the day there's just one r in the distance formula and not two separate ones.

MR. WEINSIER: That's a good question. I don't know anyone who has tackled that yet. I'll ask around. Ellen, do you want to add anything on that?

MS. COOPER: I've never done any work on the distance method. I think you really need to understand the distance between both, no matter what the guaranteed benefit is, because your discounting is using the short rates and your fund growth is using your equity. I would do a couple of different things. The first thing I would do is look at how different my answer is when I just use my S&P 500 distance versus using my short rates and figure out whether I am getting significantly different answers. If I'm not, maybe I'm done, because there is a correlation. If I am, I think I would build some kind of a function where I'm looking at the distance between both. I would look at the distance between my short rate. I would look at the distance for my S&P 500, and then I would build in some kind of correlation factor. That's just off the top of my head.

MR. ROBBINS: It certainly adds some complexity to the methodology. That's not the first time I've seen a problem with a good theoretical answer given with simplistic conditions, and then actuaries have to try to apply it to a difficult actual situation. The theory might go out the window to some extent.

MR. JEFFREY S. ROTH: Is there any rule of thumb for how many scenarios you use for the representative sampling?

MR. WEINSIER: You mean in terms of the 1,000 we took down to 50? I would say that it's probably not worth doing if you're not going to reduce scenarios 90 percent. Certainly if you start with the 10,000, you want to work to get that to 500. I would think 10,000 to 500. I think that 90 to 95 percent is a good rule of thumb. That's certainly what I've seen in our projects.

MR. JEFFERY A. FITCH: On this distance formula, let's say you've generated 1,000 interest rate scenarios, but you have entire yield curves that you're generating. It seems like this methodology requires you to pick one point on the yield curve that you're dealing with. How do you then take into account the different shapes of the yield curve?

MR. WEINSIER: The formula is tricky when you don't have a minute to study it, but it does take into account the entire yield curve. The t from one to 30 is your 30-year yield curve. For duration one, you take one over 1+i for duration. That's just by itself. The second piece is going to be 1 over 1+8, that's your scenario B squared. Now you're at year two. You're going to multiply 1 over 1 over eight times 1 over 1 over seven, minus 1 over 1 over 9 and so on and so forth. When you get out to year 30, you're going to take a multiple of 30 different numbers here for each of these two pieces.

MR. FITCH: But they're all one-year interest rates in that example. What if you have an entire yield curve where you're generating a short rate and a long rate and doing different shapes?

MR. WEINSIER: We usually take the one-year rate. There's no magic answer to that. Yvonne Chueh and I have had this exact discussion, because this question comes up every time. The short-term rate usually works better. We usually use a one-year rate, but there is no rule that says you can't use a 90-day or can't use a five-year. It's hard to say. We've always used a one-year rate.

MR. GARY L. FALDE: I had almost the exact same question. We did some experimenting with the significance method for cash-flow testing just to satisfy under the new Actuarial Opinion and Memorandum Regulation (AOMR) some idea of moderately adverse conditions, so we were more focused on general account risk, but of course, existing assets and re-investment strategies are farther out on the yield curve and this was a series of one-year rates. We weren't sure whether we were picking the right scenarios to be catching the risk in our particular investment strategy. I wondered if you had applied this in a more general account sense. I think you've sort of answered the question, but do you have any further comment on experience with this? I thought the method might break down if you used five-year rates or 10-year rates in there. It's sort of a PV number and it looks like it would not make a lot of theoretical sense if you use something other than a series of short rates.

MR. WEINSIER: I think the longer you go out, the less theoretical sense it makes, based on the formula and the PV behind it. I would have to test it. I don't have any practical experience using anything other than a short-term rate. Did you try something besides the one-year?

MR. FALDE: No, we didn't, because it didn't seem like it would make sense. It was a helpful exercise to do, and we're going to think about it some more this year and try to make it fit, but we didn't use anything other than one-year rates. But you are saying on the equity side, the 1+i is the equity return?

MR. WEINSIER: Yes.

MR. FALDE: Do you do that monthly, quarterly or annually?

MR. WEINSIER: Monthly. We got 360 rates—30 years times 12.

MS. COOPER: I have two questions, an easy one and maybe a little harder one. My first question is that I thought that the Academy was requiring us for C-3 Phase II to calculate 1,000 stochastic scenarios. I'm wondering, David, whether or not you've had any discussions at all with anybody on the task force to know whether or not a representative scenario method would be acceptable if you're actually using fewer than 1,000 to calculate?

My second question is around whether or not the distance method breaks down when you have to incorporate the impact of a hedge. We've given some thought to this in our shop, and the concern is that if you are fully hedged, then the most extreme economic scenarios are not necessarily going to be your worst scenarios, because you have a hedge in place. Has there been any work, or do you have any sense of what would happen?

MR. WEINSIER: Those are two very good questions. The first one I know the answer to because we did analyze that exact question. The question was: In the C-3 Phase II proposal, are you allowed to use representative scenarios? Our consensus was yes. There's a paragraph in there. I don't know that the term "representative" is used specifically, but it was clear to us that you are allowed to use a smaller sample of scenarios to represent the set the Academy recommends.

The second question regarding the hedging is a good one because the distance method based on one extreme is a good scenario; the other extreme is a bad scenario. But if you have everything hedged, who knows what's good and what's bad? So that could cause the methodology to break down. It simply would require further testing to make sure your representative set produces very similar results to your full set. If you're going to implement this methodology, obviously you're going to have to test it out, just like we did on the screen. You have to run your seriatim with your 1,000 scenarios and then run with your 50 and compare the results. Once you get comfortable with those results, then you can expand your work using that specifically.

MR. ROBBINS: I was thinking that same thing regarding Gary's question earlier about the fixed account. My thinking went the same way. If you had, say, a seven-year non-surrenderable guaranteed interest certificate (GIC) and you could buy seven-year no-coupon bonds to hedge it, who cares what happens to interest rates for the first seven years? It would only matter what happens when you have to re-invest when some policyholders renew and others don't. It's kind of the same issue. If you're exactly matched, I agree that the whole question breaks down a little bit.

Chart 1

Choleski Factorization

Correlation Matrix

	SP500	nasdaq	Russell2000	EAFE
SP500	1.0000	0.8100	0.7100	0.7800
nasdaq	0.8100	1.0000	0.7500	0.6400
Russell2000	0.7100	0.7500	1.0000	0.6900
EAFE	0.7800	0.6400	0.6900	1.0000
	0.0000	0.0000	0.0000	0.0000
	0.0000	0.0000	0.0000	0.0000

Choleski Matrix

	SP500	nasdaq	Russell2000	EAFE
SP500	1.0000	0.0000	0.0000	0.0000
nasdaq	0.7100	0.7042	0.0000	0.0000
Russell2000	0.8100	0.2484	0.5312	0.0000
EAFE	0.0400	-0.1539	-0.0831	0.9838

8 T

Chart 2

Implied Volatility – Term Structure

The square of volatility is linear in time

Forward volatility σ^2(1 period) =

σ^2(spot vol)T - σ^2(spot vol prior period)(T-1)

Term Structure of Volatility

	S&P		NASDAQ		Russell	
Maturity	Spot Vols	Fwd Vols	Spot Vols	Fwd Vols	Spot Vols	Fwd Vols
0.25	17.8%	17.80%	26.2%	26.20%	21.4%	21.40%
1	17.8%	17.80%	26.2%	26.20%	21.4%	21.40%
2	17.8%	17.80%	26.2%	26.20%	21.4%	21.40%
3	18.1%	18.69%	26.7%	27.67%	21.5%	21.70%
4	18.3%	18.89%	27.1%	28.27%	21.5%	21.50%
5	18.5%	19.28%	27.4%	28.57%	21.6%	22.00%
6		18.84%		28.43%		21.52%
7		18.41%		28.30%		21.04%
8		17.97%		28.17%		20.56%
9		17.54%		28.03%		20.08%
10		17.10%		27.90%		19.60%
15		17.10%		27.90%		19.60%

10 T

Chart 3

Risk Neutral Output: Average Annual Returns

Annual Rates:Geometric Average

	1 Yr	7 Yr	30 Yr	SP500	Nasdaq	Russell2000	EAFE	MM	Bond
Base Line:	5.62	5.99	6.64	3.49	0.20	3.27	3.63	5.62	5.87
Int Up:	5.72	6.09	6.74	3.59	0.30	3.38	3.73	5.72	5.97
Int Down:	5.52	5.89	6.54	3.38	0.10	3.17	3.52	5.52	5.76
Volatility Up:	5.62	5.99	6.64	3.26	-0.14	3.04	3.41	5.62	5.78
Volatility Down:	5.62	5.99	6.64	3.70	0.52	3.50	3.83	5.62	5.94
Shock Up S&P:	5.62	5.99	6.64	3.52	0.20	3.27	3.65	5.62	5.87
Shock Up Nasdaq:	5.62	5.99	6.64	3.49	0.23	3.27	3.63	5.62	5.87
Shock Up Russell:	5.62	5.99	6.64	3.49	0.20	3.31	3.63	5.62	5.87
Shock Down S&P:	5.62	5.99	6.64	3.45	0.20	3.27	3.61	5.62	5.87
Shock Down Nasdaq:	5.62	5.99	6.64	3.49	0.16	3.27	3.63	5.62	5.87
Shock Down Russell:	5.62	5.99	6.64	3.49	0.20	3.24	3.62	5.62	5.87

12 T

Chart 4

Risk Neutral Returns – Review the Distribution of Returns

	Year									
Dist	1	2	3	4	5	6	7	8	9	10
0.1%	(14.86)	(14.79)	(18.71)	(16.81)	(15.11)	(20.40)	(15.80)	(16.36)	(16.28)	(12.50)
5%	(7.79)	(7.63)	(8.17)	(8.68)	(8.46)	(8.15)	(8.06)	(8.30)	(8.08)	(7.27)
10%	(6.36)	(5.87)	(6.60)	(6.59)	(6.65)	(6.35)	(6.04)	(6.30)	(6.37)	(5.63)
20%	(4.30)	(3.76)	(4.28)	(4.45)	(4.63)	(4.12)	(3.88)	(4.18)	(4.21)	(3.65)
30%	(2.92)	(2.39)	(2.83)	(2.61)	(2.60)	(2.40)	(2.34)	(2.46)	(2.64)	(2.22)
40%	(1.75)	(1.01)	(1.39)	(1.25)	(1.02)	(0.87)	(0.89)	(0.98)	(1.11)	(1.01)
50%	(0.30)	0.25	0.05	0.10	0.35	0.66	0.51	0.28	0.19	0.24
60%	0.98	1.37	1.38	1.55	1.87	1.99	1.85	1.52	1.62	1.61
70%	2.73	2.88	2.95	2.97	3.37	3.46	3.24	3.02	3.08	2.95
80%	4.40	4.51	4.55	4.73	4.96	5.03	4.91	4.73	4.73	4.60
90%	7.16	7.07	7.37	7.17	7.61	7.40	7.42	7.30	7.11	6.73
95%	8.81	9.63	9.19	9.57	9.97	9.12	9.39	9.75	8.88	8.50
100%	17.01	22.26	18.53	17.76	19.63	21.85	21.76	20.15	17.41	16.50
stdevp	5.15	5.12	5.41	5.42	5.62	5.42	5.31	5.37	5.14	4.83
ave	0.05	0.42	0.18	0.23	0.41	0.55	0.58	0.46	0.31	0.43
Min	(14.86)	(14.79)	(18.71)	(16.81)	(15.11)	(20.40)	(15.80)	(16.36)	(16.28)	(12.50)
Max	17.01	22.26	18.53	17.76	19.63	21.85	21.76	20.15	17.41	16.50

13 T

Chart 5

GMWB Case Study Assumptions

- Benefit base: Initial premium
- Benefit cost: 35 bps
- Maximum annual benefit: 7% of benefit base
- Death benefit on base contract: Return of premium
- Reset election for benefit base every 5 years to Account Value
- Death benefit is reduced dollar for dollar for each withdrawal
- M&E of base contract is equal to 135 bps
- Age bands: 45/55/65/75
- Male: 50%

GMWB Case Study - Asset Allocation

	Large Cap	Small Cap	Nasdaq	International	**Total Equity**	Bond	Money Market	**Total Fixed Income**	**Weighted Avg**
L	6%	3%	2%	1%	**12%**	29%	59%	**88%**	27%
M	27%	11%	10%	7%	**55%**	25%	20%	**45%**	28%
H	37%	14%	22%	6%	**79%**	6%	15%	**21%**	45%
Total									100.00%

T

15

Chart 6

GMWB Risk Profile – PV Claims and PV Premiums

GMWB Case Study

Distribution of PV WB Claims and PV WB Claims -PV WB Premium
Assumes 100% election of maximum 7% each year

	PV Claims	PV Claims - PV Premium
Mean	(1,221,619)	21,852,558
Standard Deviation	4,075,855	7,530,020
Skewness	(5)	(1)
Minimum	(26,739,514)	(9,793,925)
Maximum	-	41,876,823
1% percentile	(24,041,112)	(9,238,347)
5% percentile	(7,322,740)	10,644,558
10% percentile	(2,562,346)	16,153,441
20% percentile	(6,394)	20,525,071
30% percentile	-	30,299,394
40% percentile	-	28,057,781
50% percentile	-	25,232,391
60% percentile	-	23,633,190
70% percentile	-	22,487,422
80% percentile	-	21,056,002
90% percentile	-	20,141,149
95% percentile	-	19,440,597
99% percentile	-	17,787,236
% below zero	23	2

T

16

Chart 7

VA Hedging Liability Greeks Time 0 Sensitivity

			Run 1 Drift no Skew 1000		Run 2 Rebalance no Skew 1000		Run 3 Rebalance with Skew 1000		Run 4 Rebalance with Skew (rate = 1.0) 1000	
Fair Values & Greeks			GMWB Claims	GMWB Prems	GMWB Claims	GMWB Prems	GMWB Claims	GMWB Prems	GMWB Claims	GMWB Prems
1	Baseline	-	6,568,635	(40,765,680)	11,340,646	(40,164,850)	13,165,176	(40,066,570)	17,627,771	(27,570,050)
2	Index Down	0.010	6,806,260	(40,504,579)	11,713,563	(39,900,063)	13,525,031	(39,804,213)	18,146,514	(27,391,201)
3	Index Up	0.010	6,332,143	(41,024,300)	10,993,371	(40,427,671)	12,786,004	(40,326,654)	17,117,903	(27,759,868)
4	Vol Down	0.010	5,451,801	(40,937,061)	9,568,421	(40,334,716)	11,216,611	(40,235,941)	15,722,490	(27,575,418)
5	Vol Up	0.010	7,852,837	(40,583,720)	13,208,178	(39,990,797)	15,154,233	(39,892,459)	19,607,391	(27,608,165)
6	Int Down	0.001	7,208,896	(40,703,710)	12,070,433	(40,120,262)	13,931,584	(40,022,657)	18,479,050	(27,536,036)
7	Int Up	0.001	5,972,965	(40,826,333)	10,645,586	(40,208,472)	12,415,585	(40,109,257)	16,806,444	(27,615,565)
8	Russell Down	0.010	6,609,893	(40,713,968)	11,404,381	(40,113,015)	13,233,675	(40,015,228)	17,729,970	(27,554,140)
9	Russell Up	0.010	6,523,121	(40,817,281)	11,270,724	(40,216,598)	13,072,081	(40,117,824)	17,524,704	(27,581,533)
	Delta - SP500		(23,705,834)	(25,986,022)	(36,009,616)	(26,380,376)	(36,951,321)	(26,122,069)	(51,430,549)	(18,433,362)
	Delta - Russell		(4,338,616)	(5,165,614)	(6,682,837)	(5,179,162)	(8,079,742)	(5,129,814)	(10,263,259)	(1,369,633)
	Gamma - SP500		11,325,675	24,816,738	256,431,610	19,667,812	(193,175,067)	22,729,744	88,744,240	(109,688,074)
	Gamma - Russell		(42,556,428)	1,116,619	(61,864,862)	879,157	(245,966,749)	876,555	(8,680,792)	44,260,405
	Vega		120,051,811	17,667,053	181,987,819	17,195,952	196,881,082	17,174,097	194,245,057	(1,637,360)
	Rho		(617,965,472)	(61,311,491)	(712,423,141)	(44,104,843)	(757,999,529)	(43,299,769)	(836,303,220)	(39,764,068)
	Convexity		44,591,247,103	1,318,437,026	34,727,738,034	967,283,349	16,815,938,517	1,226,021,062	29,951,703,090	(11,501,498,578)

Time 0 Liability Cost: Run 1 = 8 bps Run 2 = 14 bps Run 3 = 16 bps Run 4 = 32bps

T

18

Chart 8

A New Approach - - - Leveraging Technology

- Stochastic simulation approach:
  - Nested stochastic technology

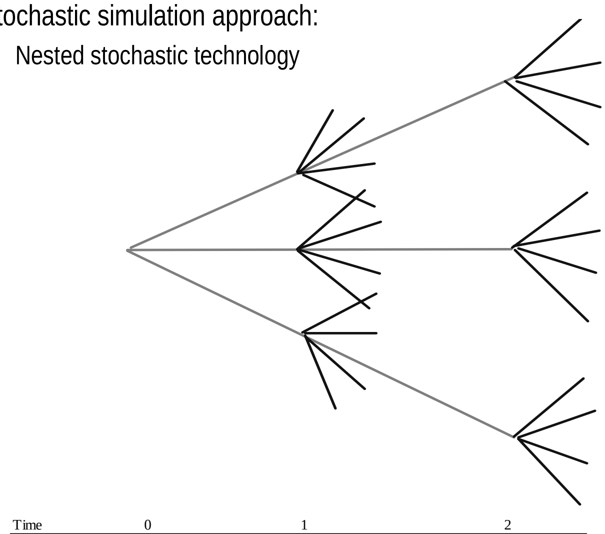

T

20

Chart 9

Pathwise Liability Greeks Under Alternative Deterministic Experience Scenarios

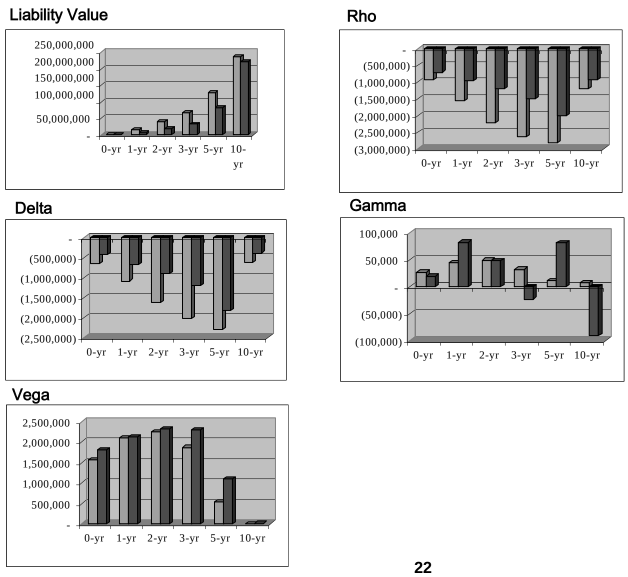

T

22

Chart 10

Central to this methodology is the notion of distance (D) between scenarios

- D is a function of the two scenarios' risk parameters

$$D = \sqrt{\sum_{t=1}^{30} \left(\prod_{k=1}^{t} \frac{1}{1+i_k} - \prod_{k=1}^{t} \frac{1}{1+i'_k} \right)^2}$$

where i_t, t = 1, 2, … , 30, is an interest rate path consisting of one-year interest rates for 30 years

- Representative scenarios are chosen such that they are "far apart" from one another

Chart 11

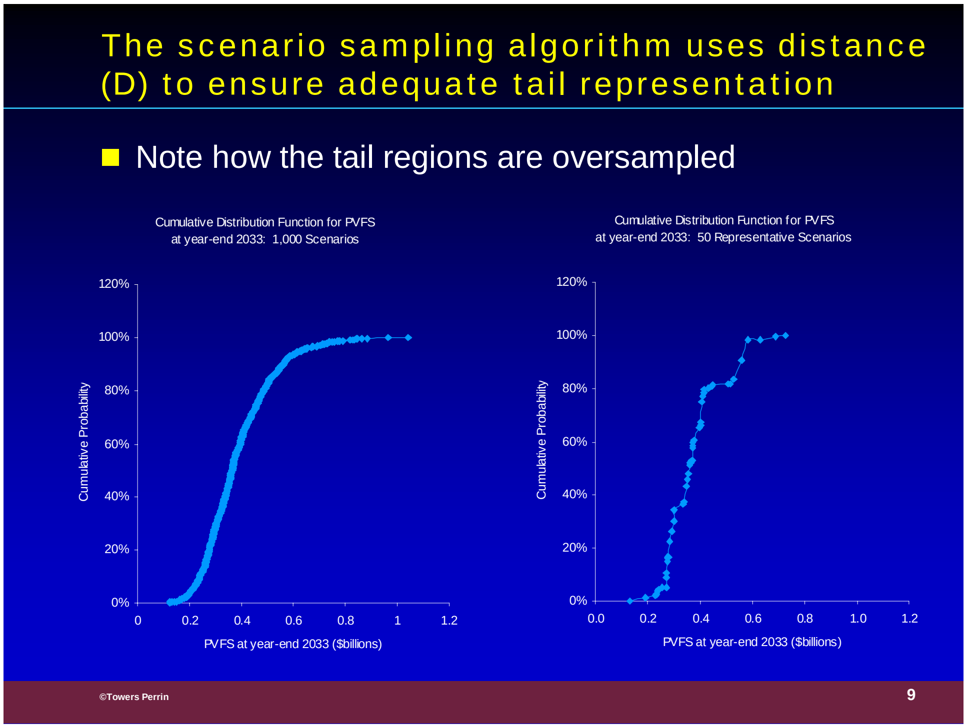

Chart 12

2. Choose the representative scenarios

- 50 (n) representative scenarios were selected using algorithms described in Chueh's paper (2002)
- Relative present value distance method

– $$D_1 = \sqrt{\sum_{t=1}^{30} \left(r_t - r'_t\right)^2 \bullet V^t}$$ (method 1, M1)

– $$D_2 = \sqrt{\sum_{t=1}^{30} \left(\prod_{k=1}^{t} \frac{1}{1+r_k} - \prod_{k=1}^{t} \frac{1}{1+r'_k} \right)^2}$$ (method 2, M2)

Significance method*

– $$S = \sqrt{\sum_{t=1}^{30} \left(\prod_{k=1}^{t} \frac{1}{1+r_k} \right)^2}$$ (method 3, M3)

*S refers to the significance of a scenario, as defined by Chueh (2002), and is used in a slightly different method in which each representative scenario has an equal probability of occurrence

©Towers Perrin 17

Chart 13

The relative present value method with distance measure D_2 gave the best results

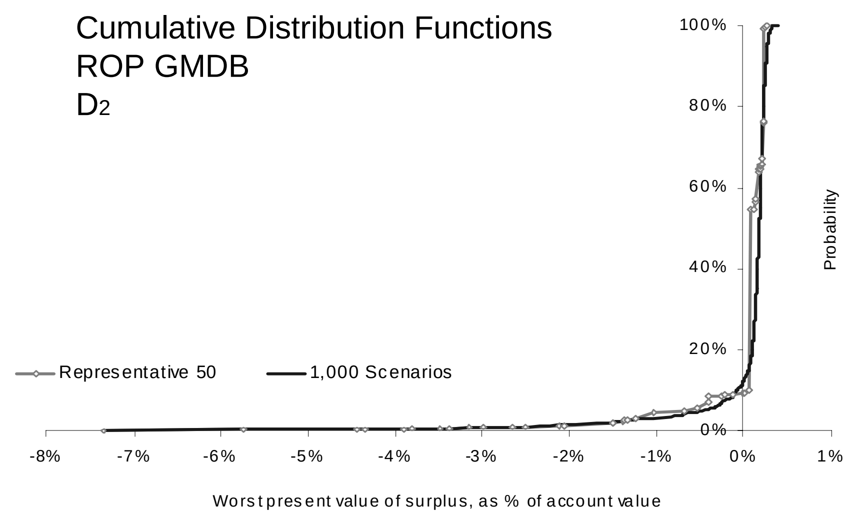

 23

Chart 14

Good fit in the tails is observed since D_2 forces representational scenarios into the tail

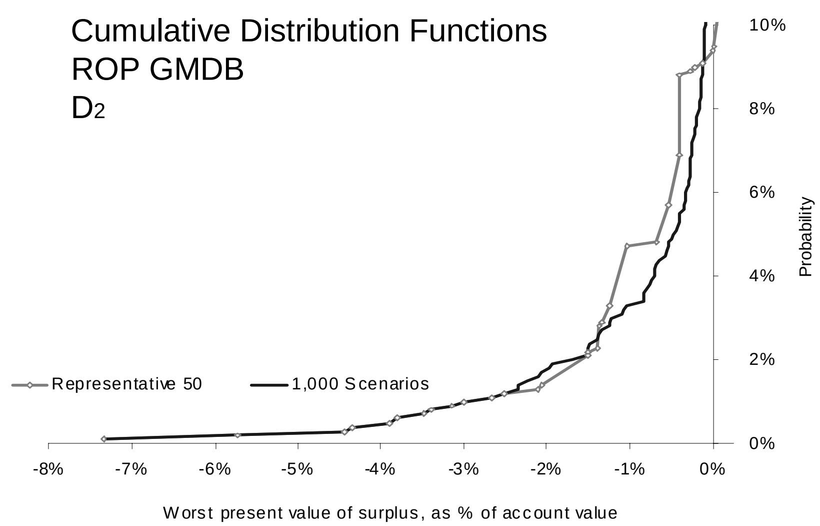

 24

Chart 15

S representational scenarios are evenly spaced
— Fit is good in the middle, but poor in the tails