

RECORD, Volume 30, No. 3*

Annual Meeting and Exhibit
New York, NY
October 24-27, 2004

Session 140PD Better Pricing in an Uncertain World

Track: Product Development

Moderator: Timothy E. Hill

Panelists: Tim E. Bennett
Timothy E. Hill

Summary: Pricing methodologies continue to evolve. We have moved from deterministic models to the use of multiple scenarios, and now we have the ability to stochastically model our risks. It turns out the question, "How much profit are we making?" has multiple answers.

MR. TIMOTHY E. HILL: I'm with Milliman, and I chose to speak on universal life (UL) secondary guarantees. I asked Tim to speak on the variable annuity (VA) guaranteed benefits because these are some of the hotter areas of pricing right now. They have some of the additional complexities that other products don't, so I thought they'd be good case studies to use for this "better pricing in an uncertain world" discussion. Both of these products will highlight a couple of major areas of uncertainty, one being the economic environment. Tim's is more on the equity side, and mine's more on the interest rate side. They also focus on regulatory uncertainty as far as reserving, capital and those kinds of issues.

Let me provide a little background on my topic. In a summary provided by the Life Insurance Marketing and Research Association (LIMRA), UL sales have been doing well the past few years, much better than variable universal life (VUL) or variable life (VL), and that was mostly due to the equity market downturns a few years ago. Term has been pretty flat. Whole life has been relatively flat, and the market in

*Copyright © 2005, Society of Actuaries

Note: All handout materials are available through the link on the table of contents.

general has been pretty flat, so the story, so to speak, has been the success of the UL markets in the past few years.

The UL premium was up by 28 percent in 2003 over 2002 and accounts for a lot of the new life sales. The reason for these new life sales is the idea of the long-term secondary guarantees that are being included in the products these days. What do I mean when I say that? I'm saying that a secondary guarantee might provide coverage to ages 90, 95, 100 or maturity, whatever the particular guarantee is providing, based on the payment of a certain premium or on a positive shadow account. A couple of different mechanisms can provide these long-term guarantees. These have been so popular because essentially accumulation UL has fallen out of favor with interest rates being so low.

People aren't buying the accumulation story as well as they have in the past. Companies are looking for older individuals with more wealth who are buying more life insurance now than they used to. So these people have a little bit of extra income. It's not just the "buy term and invest the rest" kind of people who are looking to buy life insurance for as little cost as possible to have a little extra money that they can put in. They want it to be there so that they can transfer money to their heirs, pay estate taxes and other kinds of reasons. The UL market's slanting more toward these guarantees, and it's also an older market than it was a few years ago.

The low-interest-rate environment has essentially driven people away from the accumulation side. Also, on the VL and the VUL side, people were not so excited about those in the past few years based on equity market volatility and the markets being down. Like I said, the market for these long-term guarantees on UL products is older, middle-to-upper-income people; a death benefit-oriented sale, not an accumulation sale. They're not looking to save money for retirement or take out loans in substantial portions or anything like that. It's almost purely a death benefit sale. You also are going to see more single pay, seven pays and 10 pays, i.e., shorter pay structures in these UL products than you would have seen several years ago, when most sales were level-pay type. Again, these are the people who have a little extra money and who would just as soon be done with paying premiums in seven years versus paying something every year.

I'm sure most of you have heard about shadow accounts, and it sounds like a shady, back alley deal or something. Really, the magic or the nice features that the shadow account can give you are that it will allow you to provide a secondary guarantee like we're talking about but still have a flexible-premium funding pattern. An early shadow account behaves a lot like the regular account value. It has costs of insurance (COIs), premium loads, per-unit loads and an interest rate. All it does is give you a higher interest rate than the contract guarantee and maybe higher than the current present rate on the contract, so it's saying that as long as you have a positive shadow account, your contract is going to remain in force regardless of whether you have positive account values or not. What that does is enable that

person to say, "I want to have coverage to age 90, and I want to 10-pay the contract." The reg illustration software can show, given this shadow account structure, the 10-pay premium that you'll pay for the next 10 years and the coverage going out to age 90. People can still change. This is still a UL contract, and there's still a flexible premium. If they decide they would like to not pay that much so they'll have a little less of a guarantee or they want to make it go all the way out to age 100, there is still that flexibility.

Prior to shadow accounts you had specified premium guarantees. Typically this meant that as long as they paid a certain cumulative premium, a guarantee was provided that the policy would remain in force even if the account value went to zero. The line between shadow accounts and specified premium guarantees is definitely becoming fuzzier. Specified premium designs have increasing premium requirements and sometimes there are interest credits, so the two are becoming similar where they used to be distinct camps. You either had a shadow account or a specified premium. Now the area in between is becoming grayer.

I'm going to run through a couple of different shadow account structures. First we'll talk about a vanilla shadow account. This is what I was describing as being kind of an early emergence. There's nothing all that flashy about it. You have COIs that are a level percent of 1980 CSO, and that's for single-segment issues having to do with reserving. They weren't that different from the account value COIs. You have loads that were not that different from the account value loads and so on. The interest rate, like I said, maybe was a little bit higher, and that's where you were able to buy this guarantee potentially a little cheaper than they could illustrate on a current basis. The only environment where this became all that different from the account value was when crediting rates would fall, and then the shadow account would kick in, so that was how people were buying the guarantee. The reserve implications of this kind of a structure are a little bit higher reserves than some of the other structures we're going to talk about, so often companies would use financial reinsurance to move some of those reserves off their books and onto a reinsurer.

More aggressive shadow accounts were created to potentially change the pattern of the shadow accounts. Again you probably have a level percent of 1980 CSO. You have some loads that apply up to the level-pay premium, but then in excess of that you might have a significantly higher load. The loads would vary often by age, sex, class and duration. Essentially you'd have a specific premium that you wanted the people to pay and would fine-tune your shadow accounts so they could pay that premium and have the guarantee. The shadow account loads aren't typically all that different from the account value charges in these types of funds. There is a base crediting rate but then some fairly substantial bonusing at times, so if they made it through the early years, they were potentially getting some substantial crediting rate bonuses in their shadow accounts.

Multiple shadow accounts introduce an extra dimension of complexity. There are several of these out there also, where maybe there are three different shadow

accounts. There are certain rules that say premiums go into bucket one, two or three based on the following: Is the shadow account balance positive? Is the premium in excess of something? Is it below something? There are all kinds of rules for determining which bucket the premium would go into.

This type of a structure helps you to address the flexibility of the premium pattern, which is a great benefit for the customer, but the company might not be necessarily wanting to have a ton of minimally funded UL or a ton of single-pay premiums, so with a three-pronged type of approach you can reward more the patterns that you're looking to have. If you want level pays or 10 pays, something like that, as long as those premiums are typically going in bucket two (let's make one up), you can provide lower charges and a better crediting interest rate in bucket two than the others and therefore encourage those kinds of sales, while discouraging some kinds of sales that you don't want.

The customer can still pay whatever he or she wants to pay, but you're pointing the customer in a certain direction. There certainly were abuses in some of the early products as far as people buying a UL with a secondary guarantee but only intending to have a five-year term. They could get a cheaper premium with this kind of structure because the minimum premium was less than the company's five-year term product because they had an old five-year term. Or companies saw instances where they hadn't looked at their single-pay products all that much when they priced the product, so all of a sudden they were giving these great guarantees at pretty low single premiums, and they weren't profitable. This kind of a structure can address some of those issues.

The last type of shadow account is per-unit charges only. There are just a few of these out there, where essentially you have no COIs. Most of the charges are per-unit charges versus net amount of risk-driven charges. It has a slightly different dynamic than some of the others.

Now that we've seen what the product structure is, I want to talk a little bit about product assumptions. What kind of assumptions go along with this kind of product design? When we see assumptions, often we see some mortality improvement being used when people are doing pricing. One percent per year for male nonsmokers is typically in the middle of the range. Maybe they use that for 15 years. For female nonsmokers it's 0.5 percent, not quite as much improvement. I'm sure most of you have followed the development of the mortality tables as they've moved from 1980 CSO to 2001 CSO and have seen these kind of relationships that you're not seeing as much improvement for female nonsmokers as you are for males. For the smokers, we cut each of those back by about 0.5 percent, so that would be 0.5 percent for males and zero improvement for females. Another trend that has been seen is that female smokers are not improving their mortality much.

The next assumption is the lapse assumption, and this is an important one because again this is a death benefit sale for a long period of time, essentially. We will see

lapses of 1.5 percent to 4 percent for an ultimate lapse pattern. If it is a single-pay product, maybe you'd see a lapse rate of only 1 percent or so. If the account value does go to zero, and essentially the shadow account is what's keeping it in force, you see lapses of zero. There's no cash value for them to take out, especially if it's completely paid up, so why would they lapse? They're going to wait until they die. The illustration reg issues and the self-support and lapse support tests can get tough for some of these products, so it's definitely something to look at early on. Don't get all the way down the road on your product development effort and then say, "We need to test those." There are definitely some instances where you're going to have some problems with those tests.

The next assumptions are investment returns. Most companies are pricing these benefits as portfolio products, so they'll have some kind of portfolio that's still in the range of 6.50 percent to the high 6's. You'll often get information from an investment department saying, "This is what we expect the portfolio to do over the next 10 years," and it's probably drifting down by 50 to 75 basis points over the next 10 years. That's what they'll use as a deterministic-type scenario for pricing, at least. If you're taking more of a stochastic approach, which we'll talk about later, this portfolio product is not the easiest thing to model. We'll often use a 10-year moving average of some Treasuries, plus a certain spread, plus maybe some additional recognition of the current Treasury level, so it's more like a weighted average of a moving average and a current Treasury rate.

For crediting rates, most of these again are just portfolio yield minus a spread. We don't see a lot of going up to competitors' rates by 20 percent and 100 percent down to competitors' rates or those kind of things that you would see with accumulation-type UL sales because the focus in these is not on interest rates and things like that. It's a death benefit sale, so it's usually pretty simple—150 or 175 basis points off the portfolio, whatever the particular number needs to be to hit the profit targets, is what companies are using. The crediting rate is an important assumption combined with the investment return because it is going to determine how often the benefit is in the money. If all of a sudden interest rates go up, crediting rates go up, and essentially this guarantee that was great when we were crediting 4.5 percent is not as powerful when we're crediting 6.5 percent. It's probably not going to end up being in the money if interest rates go up substantially. The relationship between the investment return and the crediting rate is an important item to capture because it does influence how often these secondary guarantees kick in because the account value has run up.

As far as the interest-rate generators that you use to model these products, in my opinion you don't need the fanciest, arbitrage-free investment generator. You need something that's going to give you a reasonable dispersion of interest rates and enough down scenarios because those are the ones that are going to hurt you. Whether you have some reversion to a mean or not is an interesting question. If you do, it definitely pushes down the cost of the secondary guarantee at least because as interest rates go up, these secondary guarantee costs go down. As long

as they're giving a nice dispersion of their interest rates, that's probably fine. I recommend the C-3 Phase I generator. I'm not sure how many of you are familiar with that. It's out on the Academy Web site. I like to use it because it's something that the clients have easy access to, and it seems that most people accept it. It's not a rigorous arbitrage-free kind of generator, though.

Let me discuss deterministic scenarios. Stress test scenarios can be valuable in testing these kinds of benefits. One example is how low the interest rates have to go before you get an 8 percent profit margin or 8 percent internal rate of return (IRR). Also, when you're talking to upper management, if you come out with 500 scenarios or 1,000 scenarios and show a distribution, that's not necessarily the easiest way to communicate results. They'll almost always ask questions like, "What if we went through the Depression?" It seems that when talking to upper management, deterministic scenarios can be as powerful as the stochastics.

Let's walk through a pricing example. We're going to do this at a fairly high level but point out most of the issues that you need to consider when you are pricing one of these benefits. Obviously, pricing usually starts with collecting assumptions. Ninety-five percent of the time you had a prior product that you're looking to enhance, so you start there. But you need to consider where this product is going to differ from your prior UL products. Again, this is a death benefit sale, so the number of lapses is probably lower. I said this is typically a more wealthy market, so there is the potential for more short pays, and it's likely to have pretty high premium persistency. We essentially ignore loan utilization when we're pricing these benefits because that's not what the sale is all about.

If you were using a mortality assumption from accumulation UL, and maybe it wasn't quite as rigorous as your term type of mortality assumption, you probably need to use something more like your term assumption. You don't have as much of an antiselection issue with this because they are still dumping in bigger premiums. It's not that you're going to get these tiny premiums and potentially get a lot of antiselection early on, but you do need to consider this more as a death benefit sale than as an accumulation UL.

X factors may or may not be an issue with this. Depending on your shadow fund structure and some other issues, it definitely isn't as important as it is over on the term side. I'm sure most of you who deal with reinsurers know that the reinsurance market has been changing. If you're using some quotes that you got a year ago, you had probably better get them updated because it seems that the prices are either going up, the terms are changing, or there are some other issues, or maybe it's not going to be there as readily as it was in the past, so don't rely on any old reinsurance quotes that you had.

The second step in the project is to say, "Now whom am I going to sell this to? Do I have a distribution channel that can get the million-dollar face amount policies, the estate planning kind of wealth transfer sales, or do I have a channel that's going to

get me \$200,000 face amount products and therefore more of a level pay type of structure?" It's difficult to be all things to all people. It's difficult to have a shadow account/UL structure that will be the best \$100,000 face amount, level pay and the best million-dollar single pay structure. You have to narrow your focus a bit and ask what your market is and to whom you want to try to sell this product. That will greatly simplify your product development structure.

You have to protect yourself from the overly competitive points. As I said, some of the early products when people single-paid those or when they did illustrations based on a single premium were competitive and not very profitable, so you need to make sure that you're looking at all those points. You need to take a look at your 10-year term premiums and say, "Am I getting a number that's less than that 10-year term premium, and therefore will I get a lot of term-type sales I don't want?" The agents will find this point if it exists, and even more so in independent channels. You have to make sure that you're comfortable with a variety of sales and all your different age and sex and rating class combinations.

Now that you've found the market that you want to target, you can get competitive information that tells you the premiums that people are having to pay for that market. You can say, "My goal for this product is to be top five in whatever list," and you can come up with a number of premiums to say, "Here's what I want them to be able to pay for a level pay guarantee for life," for example. Then you can start to tune your shadow account structure and ask, "How do I arrive at that level pay premium so they can have the coverage that they want for life?" There also are a number of dials to change. You saw all the different flexible structures that I just went through, so you can pick one of those or use a little creativity and do something a little different. It's definitely a wide-open area. Obviously certain structures are going to be better than others, but there is still a lot of room for innovation in these shadow account structures. It's often nice to have only a couple of moving parts, a couple of pieces that you're solving for to hit the right level pay premium or 10-pay premium, rather than moving and trying to solve for everything at once. That just makes the shadow account design process difficult.

The third step is to use the assumptions that we set earlier, the premiums that we talked about before and the shadow accounts to check to see whether you have a profitable product. The odds are you probably don't, so you'll have to go back. Often this is where some deterministic scenarios are used, best estimates on where your portfolio rate's going to go. The profitability can vary a lot by cell and premium pattern, though, so make sure you're checking a lot of different types of premium patterns and all your different cells. If your profitability target is not hit, you can say, "Let me address some of my account value mechanics, not my shadow account. What if I increase my spread there? What if I have some more premium loads there? What if I raised my per-thousand loads?" Those can help, but they're often not as effective as they were with your accumulation products because death benefit sales have low lapses, and so the cash value is not usually going to be the primary payout from these products as it was with your accumulation UL. You might

increase your spreads by 50 basis points and pick up only 100 basis points of IRR, which was not at all what you saw when you were pricing it as an accumulation UL.

If you're still not hitting your profitability, it's time to revisit some of your premiums or your shadow account structure. Maybe you can't be as competitive as you wanted to be because you have high expenses or your mortality experience is not so great, so you need to go back and try another cycle. This can often take many, many cycles to get to a profitable product.

This is probably the right time to do some stochastic scenarios, to say, "I do have this secondary guarantee. How much risk am I taking on? If I model out 100 interest rate scenarios, how many of them are going to have my account value running out, the shadow account kicking in and how many years prior to maturity is that going to happen?" One measure to look at is the frequency with which the secondary guarantee becomes in the money. What I mean by that is the account value runs out, and the secondary guarantee is carrying the product.

A second measure is the value of the death benefits that you pay while the secondary guarantee is in the money. If your guarantee runs out five years prior to maturity, what are the death benefits in those last five years on a present value basis compared to your profit margin, or something like that? Is that a reasonable amount to give up in profit to fund that secondary guarantee? There also is the impact of some of these account value parameters on the value of the secondary guarantee. If I increase my spread by 50 basis points, it picks up a little profitability, but now my secondary guarantee is a lot riskier because I'm dragging down my account value by more, so it's in the money a lot more. As a result, the secondary guarantee is not going to look as attractive when I take this up to corporate and try to show them that we can handle selling this product.

It's also a good time to perform sensitivity tests. Nothing fancy is needed. Mortality and lapses, like I said, are important. Test your lapse assumption and make sure that if lapses go down to 1 percent or 0.5 percent, you're still going to be in at least a reasonable profit range. Also test your premium age distribution. You might be intending to get many 65 year olds who are going to pay level premium with a \$1 million face, but what if you don't? What if you get a distribution that's different? There's a lot of subsidizing going on in these kind of products, so you have to make sure that you're comfortable with the performance for each of the cells, but you're probably never going to get to the point that they're all earning a 13 percent IRR or something like that, so that's probably an unrealistic expectation.

Let's say you've gotten through all these steps, and you say, "I hit my profit targets, have a great product and am all excited." Now you look out there and say, "What's all this stuff about AXXX and people thinking that maybe this isn't going to be an adequate reserve, and what are the regulators going to do with all this stuff?" This is becoming an important issue with these products. Just because you are following the letter of the AXXX law doesn't necessarily mean that you have an

adequate reserve—at least that's where it seems that most companies are leaning now. They're saying, "AXXX letter of the law says here's my reserve. But in fact if I do cash-flow testing or something like that, I would have to hold something higher." Obviously if you sell this, it will be subject to cash-flow testing, so you would have to be holding the higher reserve potentially.

This question of what an adequate reserve is for this kind of a product is not easy to answer. There's a lot of flexibility in the product. People can choose to do a lot of different things, so there are a couple of different issues to consider. For those of you who are not familiar with formulaic AXXX, the idea is that you compare the shadow account balance to the premium that would be required to fund the shadow account to maturity, and that gives you a ratio. Then you take the weighted average of an annual premium kind of reserve and a net single premium reserve based on this ratio, and that's a one-minute description of a complicated process. It's all laid out in steps, but it's obviously much more than what I described. The key, though, is that different shadow accounts can lead to different reserves, all else being equal. Even with the same premium pattern, same mortality and same everything else, a different shadow account structure can give you a different reserve than other shadow account type structures.

Why doesn't AXXX give the same reserve for everybody? It's not an easy calculation to try to come up with the appropriate reserve when there is so much flexibility in these products, when they can pay whatever they want whenever they want and have a variety of different guarantees going along with those premium patterns.

The Life and Health Actuarial Task Force (LHATF) is considering the AXXX issue. It's definitely on the regulators' radar screen as far as what people are really holding for AXXX reserves for these products. Is it being interpreted too literally or liberally? That's one of the issues we talked about. Is the letter of the law interpretation of AXXX reasonable, or should you be doing something more in the spirit? What are the alternatives to AXXX? If the regulators come out with BXXX, is it going to apply retrospectively, is it going to be prospective, and what are the issues that go along with that?

There is also an Academy group looking at this. Its starting point for trying to address some of these issues is to do something similar to what's going on in the VA side. That would involve more of an aggregate reserve, stochastic modeling and a "worst present value of future surplus" kind of an approach with a 65 conditional tail expectation (CTE) kind of reserve measure. I believe its stance is still to keep investigating that. Hopefully that kind of a method will work, but I know that its also considering some other approaches.

There is uncertainty about what is going to happen to these products as far as their reserving. If you decide that you need to have something other than just AXXX, if you want to have some kind of an economic reserve or whatever label you want to put around it, there are a couple of alternatives that you can consider. One would

be to go to a gross premium reserve, where you're taking the present value of future benefits, and people do this in different ways. Maybe they're only looking at death benefits, or maybe they're also looking at lapses and cash values plus the present value of the expenses, minus the present value of premiums. Obviously there are a number of different assumptions that go into this calculation. For example, for the discount rate, do you can take your current net investment earnings rate, or do you reduce that for a little extra conservatism? The mortality assumptions are how much conservatism do you put into that? There's also reinsurance—do you incorporate that into this gross premium reserve calculation? If you do incorporate it, do you assume that you have the same rates going forward, or do you assume that you have to recapture some business or that rates might go up? Just saying gross premium reserves can mean a lot of different things.

Cash-flow testing is another possibility. Once you come up with your product, you could do some cash-flow testing just as if you were cash flow testing business already on your books. Obviously there are a lot of assumptions that go along with that. You have the interest rates. Dynamic lapses reinsurance again is an issue, and there are various issues that go along with cash-flow testing. Given that you do generate, let's say, 100 scenarios of interest rates, what's the right point to have for the reserves? Is it the 85th percentile? Is it something higher? Is it a 65 CTE? There are a number of different choices to use for cash-flow testing.

Don't just do this at issue; don't assume you sell \$500 million face of this and then do one block of cash-flow testing. March that in-force block forward five years, 10 years or 15 years to recash flow test it. Even if it's adequate today, see whether there is some future point at which it is not going to be adequate. If so, perhaps that should be incorporated into my pricing somehow.

Here's a little more on cash-flow testing. It's an aggregate reserve, so when it comes to pricing, if the reserve is inadequate, it's tough to allocate the reserve back down to the cells so you can see where you have the problems. You could cash flow test individual cells, but obviously that's a time-consuming effort. You do have diversification when you have more than one cell, so you don't necessarily want to penalize yourself that severely. You could push this aggregate reserve down to the cells using a gross premium reserve or some other kind of a method that's fairly easy to calculate that gets you to the same aggregate reserve but also gives you a reasonable allocation down to the cells. That might enable you to say, "It's my 75 year olds who are causing all these problems, so let's not sell at that range, or let's not sell at such a small face amount," or whatever the particular change might be that would address some of the excess reserves that you're having.

The reserves in excess of the AXXX are generally considered nontax-deductible, similar to cash-flow-testing reserves. If you have to strengthen your reserves to the cash-flow testing, that's not a tax-deductible reserve. The cutting edge in product design is to ask, "How can I maximize my tax deduction but minimize the reserves that I have to hold but still be adequate?" These three pieces meet in one perfect

reserve structure, so to speak, which is obviously hard to do. Now you've gone through all this, and you've said "Yes, I do need to hold a higher reserve. I'm going to use the gross premium method." Circle back again and see if you're still profitable and start this whole thing all over.

To summarize at least the fixed UL piece of this, UL is already a difficult product to price—there are a lot of moving pieces. Shadow accounts allow you to be flexible with guarantees, and they do have a lot of great features, but it's definitely an extra dimension of complexity. The regulatory issues that we're talking about are another area of complexity. The reserve adequacy is just additional layer of complexity, so this is becoming a complex product as far as the product development cycle goes. Don't even get me started on second to die because there's another layer of complexity, so those are extremely difficult to price in this kind of a structure.

MR. TIM E. BENNETT: I'm an employee of a large international company (assistant actuary with AEGON USA Inc.) with various divisions. We've been through a lot of mergers and acquisitions, so through all that I've had an opportunity to see how different divisions have handled their VA pricing. We've learned from each other along the way, and we can see how each other's methods and assumptions have grown. I'm going to share with you a little bit about how pricing has evolved. Some of it's our company, some of it is what we know other companies have done. Along the way I'm going to share some results of an independent consultant survey we had done recently. It surveyed 18 of the top 25 VA carriers, and it was interesting to see what other companies are doing as well today. I'm going to wrap it up with some discussion about some things we're doing to set up some processes for equity risk pricing.

VA pricing yesterday in general, and by yesterday I mean anywhere from the early 1990s to even a year or two ago, used a simpler model and assumptions. VAs were often priced like fixed annuities using a fixed-annuity pricing model but maybe substituting interest rates for some type of equity growth component in pretty simple models. The new business distribution of policyholders was simple.

You might have had four, five or six ages. They all might have had the same average policy size, and you assumed that the asset allocation was the same for all ages. Everybody was 70 percent equity and 30 percent bonds, for example. The equity growth component was often a deterministic level upward scenario, often upward of 12 percent or higher for equity returns. The lapses and partial withdrawals assumed in the models were not dynamic in any way, just a constant assumption, and there wasn't much capital or target surplus being allocated in the pricing model as well as reserves. A lot of times no AXXX reserves were being priced in for guaranteed minimum death benefits (GMDBs). Last, hedging was often not considered by companies. The reinsurance market was more available, and you could get reinsurance at a more reasonable cost, so a lot of times you either got reinsurance or went naked on the benefits. The markets were riding high, so there wasn't a lot of pressure to make new changes.

Then we move to the bear market of 2000-02, which forced a lot of changes in companies. A lot of companies had to write off quite a bit of DAC. They faced large increases in reserves and capital and increased guaranteed benefit claims, usually GMDB claims. That period of time forced a lot of companies to relook at their pricing. A lot of them came out with different products or pulled products and made a lot of changes in their product lines because of that.

The VA market looks as it does today probably because the bear market customers were being driven toward products that give them some sort of protection for their retirement assets. Because of that, the living benefits became popular, and today either you have them or you don't. You need a strong living benefit to generate new sales. As evidence of this, about 79 percent of new VA sales in the second quarter of 2004 were in products that had a living benefit feature.

There are three basic living benefit types, and I'm not going to go into too much detail on these, but for those of you who are not familiar, I'll give a brief definition. There's a guaranteed minimum withdrawal benefit (GMWB), a guaranteed minimum income benefit (GMIB) and a guaranteed minimum accumulation benefit (GMAB). The GMWB basically guarantees a minimum level of annual withdrawals, usually until the total initial premium is returned. For example, the most common design out there guarantees that you can take up to 7 percent of your initial premium each year until you get your premium back, which is a little over 14 years. The GMWB is by far the most popular design out there. It's in 80 percent of the top 25 contracts according to VARDS for the second quarter of 2004.

The second design is a GMIB. This guarantees annual annuity income if you annuitize, usually after a waiting period of seven to 10 years, but you have to annuitize to get that benefit. This drove sales growth quite a bit a couple years ago, and the products that still have this benefit are still driving significant assets under those benefits, but a lot of them are starting to be replaced with the GMWB, which has become more popular. I think the clients and the brokers like the GMWB because you don't have to annuitize; you're not locking yourself into anything. The third benefit type is the GMAB. This guarantees a minimum account value after a waiting period of five to 10 years has been met. Usually it's guaranteeing a percentage of your premium, such as 100 percent. It's the least popular of the three, but I've heard it's starting to gather a little more momentum in the marketplace. Right now it's only in six of the top 25 contracts.

Because of these new living benefits, there's a lot more uncertainty in pricing. The market had seen a lot of GMDBs for quite a while, and at least with death benefits there's the mortality aspect that we can roughly predict with some accuracy. But when you throw living benefits into the picture, you're talking about decisions that policyholders make while they're still living, and of course, there's a lot more uncertainty in that, so it adds an additional risk that we weren't used to before. To top that off, reinsurance is typically unavailable, or if it is available, it's prohibitively

expensive. You have to pay a lot more than you could probably get away with charging in the marketplace. All this leads to the fact that we need to get more sophisticated in our pricing to analyze and manage the risks that we're bringing on the books with VAs. We need to look at our modeling techniques and assumptions. We need to consider hedging and capital markets type analysis.

Now I'm going to get into a few of the key assumptions in VA pricing and share with you some of the issues and what other companies seem to be doing according to the independent consultant survey. The first issue is scenarios. Stochastic pricing is vital now. Most companies are not using single deterministic scenarios; they're using some type of stochastically generated equity market scenarios. This is important because the equity market growth affects so many things. It affects how well or if you recover your acquisition costs, your commissions or your expenses. It affects what type of guaranteed benefit claims you have on these living benefits and death benefits. The equity market also can affect customer behavior, which, in turn, has an effect on the company's profitability. Another thing that makes stochastic rates rather than just deterministic rates that use sensitivity tests important is that it matters how you got to a point. The sequential path of the equity market matters. For example, a couple of bad years followed by good years might be worse than a few good years followed by bad years.

With stochastic pricing there are a lot of questions that need to be answered. What type of scenarios do you use? Do you look at risk-neutral scenarios, or do you base them on historical? What I mean by risk-neutral is using risk-free rates as an average growth rate. By historical I mean using average historical equity returns for your scenario generator. About five of the 18 companies surveyed used historical, seven used risk-neutral, and six used a combination. It's important to look at both types of scenarios. They tell you different things. Historical is going to tell you what the real-world expectation might look like. Risk-neutral gives you more of a capital markets view. I recommend using both, and that's what we do in my company.

Another question is what type of scenario generator do you use? The two common ones out there are the lognormal scenario generator or the regime-switching. Almost two-thirds of the companies used lognormal, and the rest used regime-switching, and there seems to be more of a tendency toward the regime-switching model for various reasons.

Another question with scenarios is what average growth rate and volatility do you use? Do you use pure historical data, which for an equity return might give you 12 percent on average, or do you add some conservatism and drop it down to 9 percent or 10 percent? There are a lot of different practices in place out there. Some companies, almost half of the 18 companies surveyed, were using pure historical data. Others were adding conservatism, and six were using just risk-neutral scenarios.

Another interesting element that stochastic pricing brings to the table is how to

communicate pricing results. We have a lot of actuaries in our senior management, so they understand this a lot more than other people might, but it's even difficult for them to move into the stochastic realm. I'm often asked what return we are making on this VA product, and it's a hard question to answer. I usually say "I don't know. What scenario are you talking about?" Even for people in tune with actuarial pricing, it's been hard to make the move into stochastic pricing and realize that there's not one answer for the return. There are a lot of answers, and they depend on a lot of variables.

Another key assumption is lapses. As I mentioned before, usually there was one lapse schedule used in pricing, sometimes with some sensitivity tests with other scales. Now we're seeing most companies, 17 out of the 18 surveyed, using some type of dynamic lapse behavior. What I mean by that is if the guaranteed benefits become in the money, lapse rates decrease. This becomes important because it hurts your left tail, your bad scenarios, because more people persist to receive the death benefit or living benefit claim, so it affects pricing. You get a little bit of an offset because you get more mortality and expense charge revenue if they stick around, but usually it's not enough to offset the extra claims you have on the guaranteed benefits.

Typically from what I've seen and from what the consultant shared with us, there are companies using some type of linear formula, a function of the in-the-money percent with maybe a threshold that dynamic lapses don't kick in until you're in the money 5 percent or 10 percent. Also they might have a cap on how much lapse reduction you would see with dynamic lapses.

An issue of uncertainty with dynamic lapses is how strong is this behavior? It might vary by company. Companies certainly have different views of it now. Not enough historical data exist to support the assumptions. We had only two or three years of a down market to gather data on, and often it's difficult for administration systems to give us good data over various points in time on this, so it's hard to get good data to support your assumptions. One thing that's generally agreed upon is that the dynamic lapse behavior is going to be a little stronger for living benefits than death benefits, mainly because policyholders have a little more control over their decisions while they're living than when they're going to die for the most part. Companies usually implement that by having a larger cap on how much the base lapse rate can be reduced with living benefits than for death benefits. For example, you might have up to an 80 percent or 90 percent reduction in your lapses with a living benefit if it's in the money versus maybe 50 percent for a death benefit.

Bennett Slide 18 should give you a sense of the impact of dynamic lapses. For 100 scenarios, I've gathered each present value of distributable earnings and ranked them and graphed them. One line assumes no dynamic lapses, just one lapse rate schedule. If I add in dynamic lapses, it deepens my tail quite a bit. The worst-case scenario almost doubles in loss. I used a one-sided formula where if the market's doing well, it's out of the money. I didn't increase lapses, but in your pricing you

should also consider whether you see that phenomenon. If I did that, my good scenarios would look a lot worse because I'd have more people leaving in the good scenarios. That's something to consider. For comparison, I put in a graph of no lapses. The dynamic lapses is somewhere between the no lapses and the no dynamic lapses.

Especially with the advent of the GMWBs, companies are now assuming the percentage of policyholders who take partial withdrawals is more of a dynamic assumption. What I mean by that is partial withdrawals increase when GMWBs become in-the-money. There are probably even fewer historical data to support this assumption than on other death benefit or living benefit lapses, just because most of the withdrawal benefits have only been out there a year or two, so there's not much historical data to gather from here.

Bennett Slide 20 is another graph to show the impact of dynamic partial withdrawal utilization. One line assumes no dynamic partial withdrawals. I think I used a constant 25 percent of customers utilized withdrawals of the 7 percent GMWBs. If you add in dynamic partial withdrawal utilization, it again deepens your tail. About 20 percent to 25 percent of the scenarios get worse, and the bad scenarios get significantly worse. For comparison, the worst case would be that everyone utilizes future insights on that.

With dynamic partial withdrawals, it's important to look at how the benefit is being marketed. If you have a GMWB that's being marketed by your marketing and sales departments as an income benefit, it's a replacement in some ways of annuitization, and you might see typically higher utilization of partial withdrawals. On the other hand, some companies might be marketing it as an account value guarantee, but you have to utilize the guarantee over 14 years, and in that case you might see lower partial withdrawal utilization. We tend to live in our pricing world and forget about what happens to the product after we're done pricing it. It's important to stay in touch with that because it could have an effect on your pricing assumptions and profitability as a company.

Another smaller thing with partial withdrawals, and some companies may have been doing this before, is a lot of companies are splitting out modeling qualified versus nonqualified money. Especially with the qualified, you have required minimum distributions you have to take, which is going to affect your utilization of partial withdrawals. It's somewhat offset by the fact that you don't have to pay DAC taxes on the qualified money, but some companies are starting to model those separately.

I'm going to move into asset allocation. As I mentioned before in past pricing models, we might have assumed that everyone was at the same asset allocation between equity, bonds, money market, etc. Now we're starting to vary that by policyholder. If you think about it, consider two policyholders where one is 100 percent equity and one is 100 percent bonds. If you average those and price them

together, you'd have an average 50/50 mix. For some of your death benefits, for example if you have a minimum death benefit, it might not show that you have any claims because of that bond component, but if you look at the two separately, you might have a lot of claims on your 100 percent equity policyholder and no claims on your bond policyholder.

It's important to separate those out, to get a feel for what your claims are going to be. We've been varying those by age. Generally we found that older policyholders tend to invest more conservatively than younger, which makes sense. They're closer to retirement or in retirement and want to preserve the money that they've accumulated. One thing to look at, if you have the ability, is how allocation varies with or without living benefits. We tend to find that customers aren't investing more aggressively if they have a living benefit or strong death benefit, which is somewhat counterintuitive. We thought they might invest with more equity because they have that floor protection. What we have generally done is split up policyholders into three or four groups. For example, we might have groups of 0 percent to 25 percent equity, 26 percent to 50 percent, and so on.

Bennett Slide 23 shows an example of what a new business distribution might look like today with that. We have four different ages, each divided into four groups. If you dig into the numbers, you can see the 45-year-old group is investing more aggressively than the 75-year-old group. This adds a few complexities to your model. One is run time. You're modeling a lot more cells than you had in the past, but one benefit of that is you have built-in sensitivity tests, as well. If you dump out your earnings for each of these cells separately you have a sensitivity test of what it looks like for an aggressive investor versus a conservative investor.

Next I'm going to talk a little bit about capital. One question that we struggled with is the current risk-based capital (RBC) allocation of target surplus, and then there's the industry working group that's been developing C-3 Phase II. It hasn't been implemented yet, but we know it's coming. The question there is which model to use in our pricing for surplus allocation, or do we want to develop our own internal model? About half the companies are already looking at C-3 Phase II even though it's not in place yet. Five are using the current RBC, three are using internal, and a few are using none. Those add up to more than 18 because a few of the companies are looking under both C-3 Phase II and current RBC.

We have a similar issue with reserves. Do we look at current Actuarial Guideline (AG) 34 and AG 39 reserves (34 is for death benefits and 39 is for living benefits), or do we use the proposed upcoming VA Commissioner's Annuity Reserve Valuation Method (CARVM)? Seven out of the 18 are using AG 34/39, seven are using VA CARVM, and a few are assuming no reserves in their pricing. It's important to look at capital and reserves in both frameworks to see what it looks like today and what it might potentially look like. The two methods are quite different. The VA CARVM and the C-3 Phase II, I think, give you a better, more accurate picture of what the right capital is than just arbitrary numbers like current RBC. The other thing that's

important is, although the exact type of credit is still being developed, if you're going to hedge these living benefits and death benefits, you're going to get some surplus and reserve credit in C-3 Phase II and VA CARVM, and that's something that isn't technically there with the current methods. Especially if you're hedging, it's important to look at the pricing in both frameworks.

Speaking of hedging, because reinsurance is unavailable and because of C-3 Phase II, which could cause companies to put up a lot more capital, especially if benefits are in the money, more companies are at least considering hedging. About two-thirds of the companies are considering it in their pricing, but as with every other assumption, hedging consideration brings up a lot of questions. One of the questions is what volatility do you assume when you're calculating a hedge cost? Do you look at historical volatility or current implied volatility? Depending on what market you're in, you could get a significant difference depending on which assumption you use.

Similar to that, do you look at current swap rates for your interest rates and your hedging model, or do you look at a long-term average? About half the companies are looking at the current swap curve. In today's market, with interest rates being so low and volatility changing, the hedge costs we found are sensitive to interest rates and volatility, so the assumptions you're using make a significant difference. We have heard of some companies that are doing a blended approach. They might start with current volatility and interest rates and then blend over five or 10 years to more of a long-term 4.5 percent or 5.5 percent assumption for interest rates. There are a lot of methods in practice, and because it's so sensitive, it's important to consider these things.

I'm not going to talk about profit targets too much because I know Tim's going to talk about them, but with stochastic pricing and the nature of these VA benefits, there are a lot of different ways to look at your profitability. You can look at a CTE major, which seems to be becoming more common, especially with C-3 Phase II using that type of major. CTE 90 would be the average of the worst 10 percent of scenarios. You can look at a straight present value of earnings or profits, and then the question becomes what discount rate do you use? Do you use a constant rate equal to your company's hurdle rate? Do you vary it? If it's a scenario where you have negative earnings in a year, do you discount that at a smaller rate closer to a risk-free rate? There are some questions about that. You can look at an average return. You can look at percentiles. There are a lot of things you could look at, and in the survey we found that companies are doing it a lot of different ways to measure the cost of guarantees and their pricing.

In summary, there are a lot of methods out there; big players in the market are using a lot of assumptions; there are a lot of uncertainties. How do you manage this? I want to share with you some things that we are doing at my company that may not give you the answer, but they give you a process to help discuss and feel more comfortable with however you decide to do it. Three things we've done at my

company are quantify existing equity risk, create equity risk working groups and develop equity pricing policies. I'm talking about just the equity risk-type concept, but it can be any type of risk. We're considering putting similar groups and policies together for interest rate risk and other risk. I'm talking about an equity framework, but you can apply it to lots of different risks that we deal with.

In quantifying equity risk, we looked at the risk across all countries, across all units and divisions and across all product lines. We have equity risk in our VAs, our VULs and our mutual funds where we're paying out a commission and recovering that through fees. In addition, we have the potential for some companies to have equity risk in their general account and to pack a product with a long tail, long-term liability or surplus. It's important to pool everything together and see what equity risks you already have on the books as a starting point. You might not have realized how much equity risk you have already, especially if you get in the habit of just looking at your division's products. You might not realize what other divisions are doing that's adding significant equity risk.

The second thing you can do is create an equity risk committee or group. Put together product actuaries, investment professionals, risk management people and derivatives people—a little of everything—and make sure they represent all divisions that do have some type of equity risks or whatever risk you're considering and get together on a regular basis. We often meet monthly or biweekly if we have a hot topic, and we discuss a lot of issues. We discuss pricing policies, hedging programs, major assumptions such as scenario generators, growth rates for the scenario generators, hedging, etc.

It's important to utilize all the expertise you might have throughout the company that you might not have been aware of. In our company I've learned a lot from the derivatives investment people about how they would hedge these benefits, and I think in turn they've learned a lot from the product actuaries about how these liabilities might work. I think both groups, through this process, are able to do their jobs more effectively and learn a lot in the process. Additionally, we've even found a few cases such as someone in our investment division who used to be a product actuary. He had great experience in a past company that he could draw from to help us product actuaries with some benefits new to us, and we wouldn't have known that person existed had it not been for this process.

Another key thing that the equity risk group can work on is formalizing a pricing policy for equity products. We found in the consultant survey that only two of the 18 companies surveyed had an explicit policy. I found that surprising. This policy could do a lot of things. I think it should in some way try to quantify what acceptable risk you should be taking in equity risk-type products, and then you could loosely outline acceptable pricing methods without getting too detailed. After this group has developed this process, another important step is to take it to senior management and get buy-in on it. If you do that through this process, this policy should somewhat represent your company's aggregate view of risk tolerance.

Let me offer some suggestions for these pricing standards. I would suggest looking at the products in a lot of different ways. It adds to your pricing time, but it's important with these types of benefits to have a lot of different tests or criteria that you have to meet. Use some type of embedded value tests where you have to be adding value to the company a certain percentage of the time, so maybe you look at the average present value of distributable earnings or something like that. There has to be some type of tail test because these benefits have significant deep tails in a bad scenario. You need some test, the CTE measure or percentile to measure the tail. Third, you need some type of capital markets test to look at your pricing in the capital markets perspective.

The benefit of the equity risk groups and the pricing policies is if you have different divisions underwriting equity risk, you can get some consistency in how the groups pricing them. You got buy-in from senior management, and as a pricing actuary with these types of benefits, it feels good to know that I've priced some things for which senior management would be okay with the risks we're bringing on the books. In some ways it enhances our relationship with marketing. I think people in marketing and sales tend to think that actuaries are too conservative in a lot of ways, so here we have a policy that we're trying to adhere to that senior management signed off on. It takes away a lot of the tension between actuaries and marketing in some organizations, or at least it shifts the blame, which is always good.

It provides a framework for decision making. It's difficult to figure out the right price for these risks and the right level of risk to bring on the books. A lot of times you run into analysis paralysis where you're analyzing it too much and never make a decision, or you're just going by gut feel that the risk profile looked and felt right, but this gives you a less arbitrary way of giving you a framework to make decisions, and it encourages companies to be disciplined in their pricing.

There are certain benefits to these programs, but on the other hand there are some challenges or downsides. One is, as I said before, this pricing policy should represent the company's aggregate risk tolerance, but that might change over time. If you put a lot more equity risk on the books, the company might not be so keen on adding too much risk in the future, as opposed to a scenario where you had little on the books. You have to often update that to reflect changes in the environment and risk tolerance. Also, it reduces flexibility. If you have to adhere to certain strict guidelines, you get less flexibility. Some people are concerned that if you outline in too much detail how to price these benefits, you could be stifling the actuaries from doing additional analysis that they might need to do. They might get in the habit of just meeting the specified criteria they have to meet, and it might stifle them from thinking outside the box.

In summary, better pricing analysis is vital with the risks and uncertainty in today's market. There's a lot of uncertainty, and you should utilize the expertise you have

in your company to develop groups, formalize a policy and get buy-in from senior management. Right or wrong, you'll have some things that you've gotten a lot of input from, and you've gotten buy-in from senior management that should enable you to feel better about the risks that you're undertaking in the market.

MR. HILL: Let's talk about general pricing issues as far as what kind of measures there are and what kind of profit targets companies are using these days. In our experience in working with a lot of different companies, it seems to us the statutory IRR is still the No. 1 target that is used at least initially for any pricing project. Depending on the product, we'll see hurdles or goals of 10 percent to 15 percent. Obviously in competitive markets those IRRs can get slim. Sometimes we will see companies that use a risk-free rate plus a risk margin of something in the ballpark of 700 or 800 basis points. Usually the statutory IRRs are based on distributable earnings, so change in capital is included in the cash flows, and interest on capital is also included. But we all know with statutory IRRs you can have root problems, so especially with the UL secondary guarantees you can get a number of years in the future that will have some negatives, so you need to figure out ways to get around those root problems. We will typically use some kind of a Becker method where we will discount back negatives at a lower rate, such as a net investment earnings rate or something more conservative than that.

Statutory profit margin, meaning present value of income divided by present value of premiums, is another common measure that we see. There are a couple of different ways to discount that. This measure often does not include the change in capital in the cash flows and does not have interest on capital. There are different ways that you can do this, but these are probably the two most common that we start with. GAAP-driven companies obviously look a lot at GAAP results. ROEs are a typical measure. Again targets are in the same range as IRRs, often 12 percent to 15 percent. There are a couple of different ways to compute an ROE, as far as whether you discount the cash flows in the equity or not. ROE patterns are often also looked at, not just the present value or the ROE at time zero, but how it emerges over time because that's what CFOs have to explain.

Tim and I talked about stochastic modeling earlier. I wonder how many times "stochastic" has been mentioned during this Annual Meeting versus 10 years ago. Ten years ago it was probably mentioned a few times, but now it seems that almost every presentation says "stochastic" at some point, so obviously we do a lot more of this.

A big reason is that we have fast computers that can do a lot more of this. It's not necessarily the case that people didn't know back then that there were risks; they just couldn't capture them as well as we can now, or at least we think we can capture them much better. The problem, though, is that all of a sudden you end up with a big block of 100 numbers and then what do you do with these numbers? It's hard to make decisions sometimes based on more information. Sometimes it's easier if all you have is one 12 percent IRR versus this big block of 100

stochastically generated IRRs.

Some common measures that we see are companies that want to achieve their hurdle rate in at least 85 percent of the scenarios. They'll price these elective benefits so that the charges are

adequate in a high percentage of the scenarios. Tim talked a little bit about the hedge costs. If you were to go out to the capital markets and completely hedge the benefit that you're writing, would you still be profitable in some manner or other? Our percentiles and our CTEs are two different ways to look at both of those pieces.

Let me talk about other potential measures that companies might be looking at. In the prior presentation in this room they talked about the value at risk (VAR), but we don't tend to see that a lot in insurance companies yet. It might come up, but I can't think of when we've done an analysis that included generating something for VAR, so that hasn't gotten engrained in the insurance industry as much as in banking, for instance, and other investment areas. There's another firm that is talking a lot about VAR and embedded value. It's combining these two items and comparing the average embedded value to a 5th percentile embedded value as a capture of risk. That certainly is a capture of risk, and obviously it's better than doing nothing or using some other tools. But with some of the benefits we're looking at, maybe the 5th percentile still isn't that bad. Maybe it's beyond that where things get bad. You could be convincing yourself that you're fine when in fact that's not the case.

It seems as though there are some other measures that must be out there that could improve on some of this decision making. Some of the questions we ask are: Should we be selling this product? Should we be buying and selling specific blocks? How can we quantify our company's risk tolerance? Economists have a tool that they use often when they talk about how a person makes decisions. It's called utility theory. You've probably heard it called a number of different things, and it's something that's easy to understand. Because it is around lunchtime, I can use this example of the diminishing marginal return of an additional hamburger. If I give you each a hamburger, you'd say, "Great, we have lunch." If I give you two hamburgers you'd say, "I'll be pretty full, but good." If I have three hamburgers, ... and it goes on and on and on. Eventually you get to the point where each additional hamburger isn't adding much happiness, and that's basically what utility theory says. The assumption that underlies utility theory is that your utility curve, the more and more of whatever it is you get, is always upward-sloping. You don't get to the point of saying, "No more hamburgers; I want zero more." You always would take that extra one, but it's adding little happiness.

Here is an example that I learned in college long before I knew I'd become an actuary. It was an example of why you buy insurance, and I thought a lot of people should see this example. Let's say this is my person's utility curve (see Hill Slide

32). Let's say that that person currently has \$150,000 in wealth; however you want to define that. The current utility measured by the left axis is about 183 units. Let's say that this person's house is worth \$100,000. The person lives in a not-so-great neighborhood, so there's a 10 percent probability that the house is going to burn down in any one-year period. If the house burns down, the person goes to a wealth of \$50,000; if it doesn't burn down, the person stays at \$150,000. If I connect those two points on the utility curve with a straight line, that will give me the expected utility for this person.

Because there's a 10 percent probability the house burns down, the little X shows the expected utility of not buying insurance. Let's consider whether that person would be willing to give up a certain amount of wealth today by paying a premium that would provide that same amount of utility. This little purple line that I just drew in is connecting me from that expected utility to my utility curve. If we look at where it hits the red line, it turns out that the difference between the current utility, the \$150,000, and where the purple line hits the red line is \$22,500, so that person would be willing to pay \$22,500 for insurance that would guarantee that level of utility. It would be at least as good as the expected level of utility, so since we all know it's a 10 percent probability of a \$100,000 loss, that's a \$10,000 expected loss. That's why we all sell insurance. People have this diminishing marginal return risk profile. That's why they buy insurance.

What's my company's utility curve? What happens if I know that today we have X amount of embedded value or whatever you want to capture, and that gives us 100 units of corporate happiness, however you want to define that. If I suppose we were to lose a certain amount or gain a certain amount of embedded value, how would that affect our corporate happiness? That would basically define our corporate utility curve. Let's say that on an annual or a quarterly basis my company generated 1,000 scenarios based on your total in-force block, capturing all kinds of risk—equity, interest rate and mortality, all that. Then I take these 1,000 scenarios, compare it to my utility curve and say here's my expected utility for these 1,000 scenarios.

From a pricing point of view, what if I then handed out these 1,000 scenarios to all my pricing actuaries and said, "You can have any kind of product as long as it increases our corporate utility." You would run those same 1,000 scenarios through your model and add it up on a scenario-by-scenario basis and determine whether you increased or decreased our expected utility.

The same kinds of decisions can be made about buying and selling blocks of business only if your criterion is increasing or decreasing your expected utility. As far as the corporate risk tolerance, you can assign a probability to whether some kind of a deal would either gain you \$100 million or lose you \$100 million. By looking at this utility function you would know exactly what the probability of success has to be in order for you to do that.

That sounds great, obviously, but people do not tend to do this, at least in the insurance industry. Why is that? First of all, it's hard to calculate what this utility curve is going to be because you have a lot of different players involved in the process. Your CEOs and CFOs have a certain risk tolerance in mind that is going to set the curve. Regulators obviously would have a different-looking curve than CEOs and CFOs. A well-diversified investor would say, "I have a pretty steep curve. I'm willing to take on some risk because I'm diversified." Then rating agencies obviously would have a different curve, so that's why it's a difficult process to perform. I think it has some real potential for being able to answer some of the increasingly difficult questions that we seem to have to answer as we get more and more scenarios and many more numbers.