RECORD, Volume 30, No. 2*

Spring Meeting, San Antonio, TX
June 14–15, 2004

Session 56 PD
Measuring and Pricing for Tail Risk

Track: Risk Management

Moderator: Henry M. McMillan

Panelists: Henry M. McMillan
Douglas L. Robbins
David M. Walczak

Summary: Tail risk, long recognized as important to insurance organizations, is being evaluated in a new light, with new approaches. Panelists review various approaches for measuring tail risk and how to price for tail risk using capital-market consistent techniques.

MR. HENRY M. MCMILLAN: I am Hank McMillan. I am the moderator and will be the first speaker in this session. There are two other speakers. Doug Robbins is from the Tillinghast organization in Atlanta. Doug is a specialist on modeling all sorts of life insurance liabilities. Dave Walczak is from Deloitte Consulting in Minneapolis. He also does a lot of work on GAAP and GAAP auditing of various insurance companies.

I'm going to begin by introducing tail risk. Then we're going to move on to Doug to give us some real-world measures and prices. Last, Dave will present theoretical measures and prices for tail risk. I'm also going to give you an advertisement for the Risk Management Task Force and its various subgroups. The folks on that have done a lot of great work. A session like this can't possibly cover all the things that would fall under the heading of "tail risk," and those people have a lot that they can bring to the table.

Let me begin by trying to define tail risk in an intuitive way. First of all, I tend to think of tails as those events that occur away from the mean and in the extreme.

We know that any point in a continuous distribution is equally likely to occur. However, your notion is that a tail event is something that is relatively unlikely. To get that far out is not something that you would expect to happen.

When we define risk, we're normally thinking about bad outcomes. If you have a tail risk that is to win the lottery, you will seldom have people look at you and say, "That's a risk? That's bad?" No, you tend to think of risks as outcomes that are costly or undesirable. But the way we normally go about doing things is that we frequently measure risks by means of a standard deviation or a volatility, which is a symmetric measure.

Mechanisms to deal with tail risks include variable annuity guarantees, medical stop-loss policies, various sorts of nonproportional insurance coverage, earthquake and flood coverage, and catastrophe coverage. Hole-in-one insurance and motion picture revenues are some other things that get outside of that.

In another life I was an economist. My partner and I did some work on the motion picture industry. He actually has written several papers on it. One of his findings in looking at motion picture revenues is that if you wanted to try to fit a distribution to it, you would have to have a distribution of infinite variance, which is a statistical outcome of the old adage in Hollywood that "nobody knows anything."

As far as hole-in-one insurance, if you play golf, one of the things that you're accustomed to is that if you make a hole in one, you're supposed to buy drinks for the entire clubhouse. Consequently, you find ways to ensure that. In my club we have a pool, and every time somebody hits a hole in one you get a free drink and every member gets charged one dollar to cover the costs. It works out very well, as long as you have your hole in one at your home course. If you have it some place else, you're going to have to do it some other way.

We're going to be focusing largely on the kinds of tails that go along with normal distributions, but you do need to think of all kinds of distributions that might fit various kinds of events with which we work. We work with normal distributions to a large extent because of their tractability. We know that normal distributions are plausible. You see them all the time in life. They are very good first approximations to a whole menu of events out there. When we start looking at them in the context of financial issues, we are probably going to be looking at non-linear payouts, things that happen that you get more money in the tail of the distribution than you get in the normal part of the distribution so that they're not proportional.

Options are an example of something that pays out in the tail. Frequently, you'll see compensation of senior executives that's going to have a very large payout if you have unusually good outcomes for performance in the company. You'll also see unusually large terminations if you have very bad performance and that kind of thing. The tails differ as you go along the continuum of payouts or as your preferences about the payouts differ. You need to get this non-linear thing for tail risk to be relevant to you because otherwise you're just dealing with means, and then it's all very simple.

When we look at financial markets, we find that the tails of the distribution tend to be larger than and more likely than you would expect from a true normal distribution; things like power laws and various other kinds of distributions will get you to that kind of a situation.

When I studied this many, many years ago, I had a graduate student working for the hedge fund. He was in the process of doing an elaborate breakout of the distribution of one day's returns on the stock market. This was around 1984, so about 20 years ago. What he came up with was that stock price movements within a day tended to have a long neck because often things didn't trade because of the minimum trade increment at that time. It had a long neck and broad shoulders because you would tend to have more one-eighth price movements than you would expect, again, because of the discreteness. You would be sort of narrow at the hip, and then you'd have big, fat tails. When you do have fat tails relative to a normal distribution, the message is that you have to have something that is narrow relative to a normal distribution, too.

To get fat tails, there are two simple ways to think about this. You're mixing distributions. Regime switching is where you go from a bull market to a bear market. This is the kind of thing that we're talking about possibly using now for risk-based capital to fit the historical equity returns. You jump from one return to another. You cause a situation where the tail of the aggregate distribution looks fatter than you would think it would be if you had one underlying distribution. The GARCH family, that is, the Generalized Auto-Regressive Conditional Heteroskedasticity family, provides a way for the variance to change continuously. You have innovations in the variance the same way that you have innovations in the mean. Again, it has that characteristic of generating fat tails relative to a normal distribution.

These kinds of things occur, if you're looking at the time series of various economic variables, typically because there's some change in the underlying economic environment. These kinds of discrete shifts can occur because of changes in the economic policy or because of the nature of the politics in the world, as well as innovations that might occur in technology and things of that sort. Also, shifts occur due to changes in the economic behavior as people learn as they go through time. The point of this is that when we're dealing with these kinds of tail risks, we have to be cognizant of the fact that our tails might be fat relative to normal distributions.

I promised to give a little advertisement for the Risk Management Task Force subgroups because they've done a lot of things. The Equity Risk Subgroup has prepared a recommended reading list, which does contain a lot of references that you might want to go take a look at. You can access this recommended reading list by going to the Society's Web site.

Another thing, with reference to the regime switching argument, that you can see very easily is in the *North American Actuarial Journal*; an article by Hardy in 2001 provides a nice review. With respect to the ARCH and GARCH family, I would encourage you to try to find *Journal Of Economic Literature* from 2003. There's an article by Poon and Granger talking about forecasting volatility in financial markets. I specifically pointed that one out because Clyde Granger is one of two econometricians at the University of California in San Diego who won the Nobel Prize. Robert Engle is the person who has been most credited with doing the work on the GARCH family. Clyde Granger worked on co-integration, which has a lot to do with issues about whether you should expect the equity performance in the next ten years to replicate the equity performance in the last ten years and if there's something going on that should drag equity prices back down to some long line trend.

There are a lot of things in the subgroups and a lot of these folks whom you could access, especially if you're looking at pricing and measuring pricing tail risks, credit risk management and economic capital. The enterprise risk management session is, perhaps, less of that, but the equity modeling and the extreme value models are useful. The extreme value models, obviously, point home for all of this. Then we also have ones on health risk management and policyholder behavior in tail. The issue there is that peoples' behaviors tend to change when you get off into the tails of distributions, which creates even more changes in your forecast. Pricing for risk, risk-based capital covariance and risk management metrics are relevant and good sources for further reading and study on these subjects.

MR. DOUGLAS L. ROBBINS: I'm going to be talking about measuring and pricing for tail risk in the real world. When I say that, I'm referring to the fact that the core syllabus that you probably read when you decided to come makes some reference to tail risk, and then it talks about measuring these things with regard to capital markets. Dave Walczak is going to get to that.

First I want to talk to you about the difference between a capital markets approach to the tails and then a real-world approach. I'm going to talk to you in five steps. First I will give an introduction of *whose* real world and the advantages of real-world scenarios. I will describe measuring risk with two real-world sets. I will discuss pricing for real-world tail risks and then I will discuss some miscellaneous issues. Finally, I will give a summary. I've approached this economically. My case studies are mostly going to be of the variable-annuity-with-guarantee sort that Hank brought up because it's just, I think, the easiest way to represent some of these points.

The view of the real world depends on where you live. Perceptions of what is realistic, or likely, will vary, depending on the source. When you're putting together an economic scenario set, the first two things you might think of are expected return and the volatility, and your rate of net expected return. Even those can be subject to contention. You can have a lot of difference, particularly in your tail, between two sets of real-world scenarios, depending on such things as the assumed distribution of rate movements. Hank alluded to that somewhat. Lognormal progression would be the closest thing to normal theory. If you take the log of the movement, you would end up with a normal distribution. That would be one theoretical approach to interest rate or equity rate movements, but certainly not the only one. It's probably one of the most aggressive if you're trying to get a conservative price for a benefit. Mean reversion is the tendency to tend toward a long-term growth rate. You can have the same mean return for any given year in your projection and the same volatility, but you can have a much different result after 10 years in the spread of your possible situations you end up in, depending on whether you believe in mean reversion or not. For equities you might assume that you're going to get 8 or 9 percent a year even though there's going to be a lot of volatility. If rates start out by doing very well, do you assume that there's now a tendency to come back toward an overall 8 percent growth? If they start out collapsing, do you assume that there's now a tendency to have rates higher than normal? If you assume that, it causes a tight spread of possible outcomes after 10 years. If you don't assume it, it causes a much more dispersed spread. That's often true even for sets with very similar expected return and volatility.

What am I really saying? You could have a pricing actuary come up with a set of real world scenarios who wants to come up with a cost for a guaranteed benefit that will sell. They come up with a reasonably conservative set for equities—expected return of 10 percent and volatility of 16 percent. For many historical periods you can justify those two parameters. But the process is simple lognormal and there is mean reversion.

Now you have $1 billion of benefit in force. The valuation actuary steps up and says, "I'm going to value this over some real-world scenarios." His expected return is 8 percent and his volatility is 18 percent because he doesn't want reserves to be inadequate. His process has jump diffusion and he assumes no mean reversion. Maybe the benefit cost comes out a lot higher and maybe there is a risk higher than it was priced for. You never know.

For the regulatory view, especially now with C3 Phase II coming out, the expected return and volatility are irrelevant. It doesn't matter what your expected return or your annual volatility is, but you have to calibrate the tail of your scenario set to very conservative specifications that are outlined in the guidance. All of these scenario sets are real world, but all of them are very different. What makes them different is the point of view of the person who's creating them.

What do we do now? Is it possible to reconcile these views in measuring and pricing for tail risks? I maintain that it actually can be necessary. In our next few sections we'll see why. I'll follow a couple of case studies on variable annuity guarantees and illustrate how this can come about.

First I want to talk to you about some advantages of real-world scenarios, though, because we're going to hear about a capital markets approach later. I want to point out why it might still be worth talking about a real-world approach and a few aspects of that. The first is that the real-world set of scenarios gives an impression of plausibility. You can say that your scenario set is calibrated from stock market returns from the 1950s to present. That creates an aura (kind of) that you have built these based on something realistic. If it's happened before, it can happen again. You may fit the data to some probability distribution. Again, this is something that your actuarial audience and, to some extent, any technical audience, will be used to. You can express expected results and volatility in terms that people are used to hearing.

Descriptions of risk-neutral scenarios are often not appealing in this way. What do I mean when I say that? Let's think about what we might be doing. Let's say we're pricing a guaranteed minimum accumulation benefit over a 10-year period for a variable annuity. That is a guarantee that if you pay $10,000 of premium now, that in 10 years, no matter what funds you invest in and no matter how the markets do, you're going to at least get that $10,000 back. That's something like a put option. If you value it over a realistic scenario set, especially one with mean reversion, you might find a very small cost after 10 years because, especially if you run a small number of scenarios like maybe a hundred, you might only have one or two scenarios that puts you below your premium at the end of the 10 years. When you sum those, divide by a hundred or however many, you're going to get almost no cost.

Risk-neutral scenarios are a capital markets-based approach. You assume that the price the markets would give you for a 10-year put option is correct, and it's going to build the scenarios based on that. The risk-free return is what it is. Therefore, the only way you can get the correct price to come out after 10 years is to jack up the volatility over that 10-year period. What you get is a whole lot of scenarios, for an obvious reason, that produce negative returns over a 10-year period, which is something that people are just not going to be used to hearing.

If you look at stock market returns over the past hundred years, even over the worst 10 you could possibly pick of the Great Depression, there is no 10-year period where that would be true. When people think real world and then you show them the results of a set of risk-neutral scenarios, it's just not going to be appealing. They're going to look at it and think that that's not possible. That couldn't happen. Well, as Dave will say, that's true, and it doesn't make any difference. One of the appealing aspects of real-world pricing is that the scenarios do seem plausible.

Another thing about realistic scenarios is that they're easy to use for multiple purposes within a run. Most economic scenario sets have the primary purpose of projecting cash shortfalls based on some set of guaranteed element. In this case, you might be pricing a death benefit or a living benefit for a variable annuity and you want to know how many scenarios produce a positive cost for these benefits. A real-world scenario set can simultaneously be used to project policyholder behaviors. Hank alluded to this.

Say you bought, again, something that is tantamount to a put option at the end of 10 years. If after seven or eight years your fund has done very poorly, you're going to be a lot less apt to lapse your annuity than you would be if it were doing well, and the scenario should reflect that. It's basically very easy to build a dynamic assumption off of a real-world scenario. Again, the scenarios seem realistic, and you can say, "Over this set of scenarios that are apt to represent what could happen, here's the number where I'm going to have an issue with my lapses."

If your scenario doesn't seem plausible, then it's not as obviously true that you can use the same scenario set as you were getting your market cost for options to also project policyholder behaviors because it just might not seem like these scenario sets are realistic enough to do that.

One final point that's very important to me as a statistician and as an actuary is that inferences on tails can be drawn directly from your results. That is because the scenario set is a real world, in the view of whoever is creating the scenario set, so each one represents a unit of probability. If we've got 1,000 scenarios, each one is .001 likely. Then you can build a set of results and use that, assuming all your other model assumptions are correct, to make probability or statistical statements about your possible results about your parameters.

For example, I project 1,000 scenarios to test the benefit, and my rider charge that I'm thinking of charging for my benefit appears to be sufficient to cover my cost in 950 scenarios. What can I say about my true probability of covering at least 90 percent if the scenario set and all my other assumptions are correct? It's very simple. You would say, "If my true probability was less than 90 percent, what would my probability then be of getting at least 950 scenarios that were incorrectly projecting that I do cover costs?" You set that up using a binomial distribution and you very quickly get to the answer. But I think this is about like a 15-standard deviation event. In other words, you can be very, very sure that you are covering your costs in 90 percent of the scenarios.

You can build statistical inferences based on the scenario set that's created in this way. Correspondingly, you can't do that with a set that's built just to generate a mean cost that replicates capital markets, because the tails don't have any meaning.

What are the corresponding disadvantages? The biggest one is that there's really nothing observable out there with which to demonstrate correctness. Remember, early on I gave you three possible real-world views. Which one is right? There's no way to tell. You can't say that this scenario set produces something in the future, therefore it is right, like you can with a risk-neutral set where you're actually calibrating to market costs for options. This only exacerbates the fact that history often provides a very small sample space. What do I mean by that?

Let's take a 14-year guaranteed minimum withdrawal benefit. Your horizon then is at least 14 years long. How many independent—key word is "independent"—14-year periods do we have to help evaluate a large cap stock fund? If you think, particularly as I do, that a lot has changed since the 1920s, then you should only be looking in the modern era. You might only have two, three or four 14-year periods that are independent of each other, which, drawing from our statistical knowledge, is a very poor sample. You want something like 30 periods to even start to talk in terms of statistics. Even though history can help us shape one-year returns pretty well, it has trouble with a 14-year period. It has trouble with questions like, does mean reversion exist or not? If the whole benefit could be reset after five years, then you're talking about more than a 14-year period, so the situation is even worse.

However, the use of a credible real-world set can still help a great deal in measuring risk. At least it can tell us, if the future is much like the past, what our tail results would look like. That can be important, particularly if we have an unhedgeable benefit, which some people say is the guaranteed minimum income benefit (GMIB). It would be crucial just to prevent us from getting overexposed to a given risk. If we can say we're comfortable as long as we're not sinking the company in more than 2 percent of our scenarios and we can come up with a real-world scenario set that we're happy with, then we can at least say something about that 2 percent, and we can agree that we're not going to sell more than X percent of our in force to have that benefit so that we're not exposed beyond that point.

Even if we're happy with our realistic set, it makes sense to sensitivity test more pessimistic views because you never know. Another thing that helps a lot with real world is that many of your benefits that you're testing—maybe guaranteed minimum death benefits (GMDBs)—are going to have a much shorter time horizon than the one I brought up earlier. You'll have a much larger historical sample.

Let's go on to measuring risks with two real-world sets. Remember we talked about our pricing real-world sets and something that the pricing actuary might come up with. As I said, I'm going to apply this to variable annuity (VA) and guaranteed benefits. If you're in more of a universal life (UL) world, just try to think along with me. I will refer in my summary section to some other applications, but I'm going to go along with the VA example here. We've got a new GMDB and new living benefit. Let's say we're thinking of costing a 5 percent roll-up GMIB. Let's talk about what kind of thoughts we might have.

We already know about the interest rate haircut. I'm not going to get into the fact that that could change, too; our effective roll-up, instead of 5 percent, might be 3 percent. I'll explain it quickly by saying that a 5 percent roll-up GMIB is a benefit where you're guaranteeing not the account balance to roll up at 5 percent interest a year, but it's a guarantee you put on equity funds so they could go up or down. At the end of the waiting period, the person can annuitize using a shadow fund that goes up at 5 percent a year. Now that annuitization is at a very poor guaranteed rate, which is where I get the haircut. You're going to have higher rates than you would expect on a current basis, which means the effective difference between your shadow fund and your fund value is reduced by the haircut. So the effective roll-up could be as low as 3 percent, which means, since we expect 8 or 9 percent growth from equities, in most of our real-world scenarios we expect no ultimate cost. We expect our fund to be bigger than our shadow fund after the haircut.

That means that our tail is our chief concern. What should our tail look like? What things do we need to think about in our basic pricing real-world scenario set? We should want scenarios that accurately reflect our anticipated variable annuity fund mix. With our chief concern being the tail, we've got to ensure that the way we model this fund accurately reflects what is going to happen in our tail given our distribution of fund scenarios. We would agree in my firm and with several clients that the best practice is to reflect at least a few asset classes (maybe three to five)—at least one equity class (maybe more), probably a bond class and a cash class—with reasonable correlations between them.

Real-world scenarios for long assets would typically be based on the long-term risk-free rate, long assets especially being equities. In an environment where the risk-free rate is high (and this goes along with Black-Scholes), we would need an option theory. The cost for a put option would be less because you expect higher growth in the long run from your equity funds. A risk premium would also be included, and that would vary by asset class with higher for equity and lower for bond.

Volatility would also vary by asset class. It would be highly correlated to the risk premium as the higher the risk, the higher the volatility. You would want a mix of volatility and risk premium that makes sense based on the correlation of your funds. You want to reflect whatever your view is on mean reversion, which we talked about. That forms our tail picture.

Should you consider anything else? Probably. Right now, per unit of in force, in other words, per $1,000 of initial premium, perhaps, we can project a cost in each of our scenarios or a set of costs. For every $1,000 of premium in a large cap equity fund, for instance, we're going to have in scenario one a spike up of benefits, a ratchet benefit. The next spike back down is going to cause a cost and then it will go back up. Those costs might be $10 per $1,000 in year three and $30 per $1,000 in year seven and so on. That's for that scenario.

What we don't have, and this can often be missed, is a projection of the units that are going to be in force when those benefits come due. That's the behavior assumption that we talked about earlier. Using a realistic set of scenarios, you should be thinking in terms of dynamic behavior. Is the fact that they're holding this valuable benefit going to affect lapsation? Is an annuitization benefit like a GMIB going to result in heavy utilization of annuitizations in year 10? What are your policyholders going to do? We need to formulate a credible assumption regarding policyholder behavior in the tail of our pricing of real-world set.

Some work has been done on this in the Stochastic Modeling Symposium in 2003. More work is being done in one of the committees that Hank mentioned. But what do we have at this point? Assuming the correctness of our assumptions regarding the stochastic scenarios and policyholder behavior, we have a distribution of variable annuity guaranteed living benefit (VAGLB) or, in my second case study, GMDB costs. What can we infer from the distribution? We can get a mean cost, maybe 10 basis points or 25 basis points if we divide by the present value of fund value. We have a confidence interval that tightens as we run more scenarios. If we want to know that the number is 25 within three basis points given two standard deviations, the typical rule is that you might have to run 100 scenarios. But if you want it within one basis point, you'd have to run closer to 900 scenarios.

We also have a picture of tail percentiles that's mathematically unbiased. Again, this is given the correctness of our basic assumptions and then if the real-world view we've taken is correct, then it's always true that Monte Carlo sampling gives you an unbiased picture of your tail. That is really important. As discussed before, we can also look at confidence intervals on our tail percentiles. If we want to know the 95th percentile within 2 percent, we can get it. We just have to run enough scenarios.

That gives us two potential bases for trying to set a price for our rider: mean costs or tail costs. Can we, or should we, use our analysis thus far to set a price? We can if we want to. We've seen people do this. It is common to pick a percentile we're happy with and set our price at that level—85th, 90th, something like that. The mean cost, if we use that, would cover us on an expected basis, but wouldn't really provide us with a margin for adverse deviation. In other words, it's a risk premium for us. What that means is, as the company, we're taking a risk by charging something and guaranteeing this benefit, so we really want to make more than our expected cost.

That approach would involve an assumption that if we cover the given percentile (the approach of setting it at a tail like 85th or 90th percentile), we'll be profitable over the whole range of possibility. Can we say that? Probably not. For one thing, we don't really have a distributional basis for where we're going to be if our actual scenario turns out worse than the given percentile that we picked. All we said is that we're going to cover 85 percent of possible outcomes. To say that we're profitable over the entire range of outcomes, we need to make an assumption that it's not going to get too bad if it goes beyond that. In reality, tails, as Hank pointed out, can be very fat when we're talking about economics. We have to also price in capital costs.

The new methodology that we're going to get under the proposed C3 Phase II guidance gives us a basis for bringing those costs into pricing. I would argue that it's a good basis anyway. The C3 Phase II is the new risk-based capital (RBC) guidance for variable annuity guarantees. Here's what it tells us to do. As a means of setting part of your VA's required capital, the C3 part, you adjust your scenario set or create a new set that calibrates to some specifics that the regulators are giving as calibration points. After one year, it's got to be at least this fat. After five years, at least this fat. It gives several percentiles that you've got to hit. This is the regulators' real world. It is very conservative in the tails. Project your VA, VAGLB, or GMDB over at least 1,000 calibrated scenarios. Look at your minimum present value of surplus at any duration and set your required capital at the average of the worst 10 percent of those results.

It's giving us a means of setting our RBC, our capital of over and above our reserves. It's called a conditional tail expectation, a CTE 90 result. It's what I was saying about where your 90th percentile is, but this is the distribution of how bad it gets when the conditional part, that tail expectation of expected value, is above it.

How do the costs compare for a realistic set and a calibrated set? I'm going to switch at this point to a combination roll-up/ratchet GMDB. Probably more of you are familiar with a 6 percent possibly roll-up GMDB where you guarantee if you die you're going to get at least interest, as if interest were 5 or 6 percent a year. It's the combination of that with a ratchet, where you'll never make less than any anniversary prior to your death.

I've looked at basis point costs first for my basic set of scenarios and then for a set calibrated to the regulators' guidelines. What does that tell us? In that particular case, the costs below the 90th percentile are lower. Is that surprising? No. There's really no reason they shouldn't be lower because you don't calibrate any of those points and even the 90th percentile calibration point is pretty mild. It's ones up above the 90th percentile that are very conservative, and those costs got quite a bit worse. The GMDB costs with those points aren't the key issue, only the surplus shortfall. But that surplus shortfall, the amount that you're in the hole, so to speak, is affected by your GMDB costs. If they're much worse under the calibrated scenarios, you're going to get much lower surplus and, therefore, higher RBC. That has got to be priced. That helps us look at our tail costs under these real-world scenarios.

Where does running those scenarios get us? Initially, we run our 1,000 scenarios as we issue the product. It gives us an initial required capital level for our VA including the added benefit. Many of you might find that that's already higher than you were expecting, if your benefits are rich enough. One thing it gets you is a higher initial capital, which is a higher strain. You can also repeat the costing analysis at future points in the real-world projection, and that gives us future required capital levels throughout all points in your scenario set. However, that involves some cost in terms of what I'll tell you in a second. It does give you the basis for pricing in your tail risk initially and at all the future projection points in your scenario. It puts you in a situation that we're starting to refer to, by the way, as stochastic-on-stochastic projections.

Now we're ready to price for real-world tail risk. I'll just tell you how I would recommend doing it, based on C3 Phase II particularly. When we price guarantee riders, what are we trying to accomplish? There's no initial commission on them so that affects with a very low strain on the rider itself. Strain comes only from capital required. Also, that strain probably builds as the revenue accrues, and the costs come very late in your projection. It's probably not a good candidate for looking at the internal rate of return (IRR) on the transaction.

If we try, on the other hand, to equate return on assets at a reasonable discount rate, like something around the earned rate of 5, 6 or 7 percent, then we're putting a reasonable cost on RBC-driven strain because the RBC-driven strain is going to only earn the after-tax rate and you're discounting at the before-tax rate. But, at the same time, it's a low enough discount rate that it puts a reasonable value on the actual cash costs that occur way out in your projection. It's a way of reasonably equating what you're going to charge for the rider with the cost both of capital and actual.

In the case study that I put together on GMDB, I created the distribution of actual costs in a base scenario. If I'm thinking of putting together this very exciting, rich rider, a roll-up/ratchet combination, and adding it to a guarantee that is only return of premium, I might find that at the 85th percentile my cost went up 25 basis points. So I might charge 25 basis points under my original. Let's price for this by just equating a percentile. Is that good enough? The mean cost increase is only 20. I do get a five-basis point margin above my expected cost as a risk premium. However, over the worst 10 percent of scenarios in my case study, the calibrated set increased my GMDB cost by 43, not 25. I'm missing 18 basis points, and that shortfall is going to increase the amount of my surplus loss in all of my RBC scenarios. In other words, I might have an initial RBC increase of about 50 or 60 basis points of my fund value. Is a five-basis points charge going to cover that? That gives you a cap factor of 10 to 12, which are not reasonable expectations for pricing. The moral of the story is that your pricing needs to capture the capital costs as well as the potential for actual. That is basically how, on a real-world basis, we would recommend pricing for tail risk.

How about pricing over a stochastic-on-stochastic run using C3 methodology not only at issue, but after one year, after two years and after three years? That will allow us to build in capital costs driven by tail scenarios not initially, but as the tail scenarios in our base run start to happen. If the capital costs shoot way up, we'll be building that in. It also allows us to pick up cash costs in the tails of our own real-world set. We might be satisfied that we're adequately covering tail costs, both capital and actual, but there's a disadvantage. The main one I'll get to in a second. The first disadvantage, though, is that although you can measure and price for tail risk, you're not actually mitigating it. Only by doing some kind of hedging program can you mitigate it. If you priced this way, you're okay for a while. But if a really bad scenario starts to happen, you'll notice that all your parameters start to change, and you'll have to redo your analysis and see where you are. A hedging program might help where simply analyzing and pricing for the risk does not.

The biggest problem with the stochastic-on-stochastic measure is computer run time. Think about it. You're running, conservatively, 500 scenarios that fit your pricing view of real world. Along with those 30 future projection points, you're going to branch out and do 1,000 C3 Phase II-required regulatory scenarios. That's 500 times 30 times 1,000, which is 15 million. It's unlikely that your computer capability is going to handle that.

What can we do? The RBC piece is based on a CTE 90 requirement, which renders 90 percent of your scenarios in every distinct run irrelevant. It's not going to be the same 90 percent in every case, but they should be highly correlated. A reasonable choice of 100 scenarios could suffice for your purposes. That cuts you down by 90 percent of your run time. Again, the pricing piece, your 500 scenarios, is based on a set of scenarios that could be capable of representation by a smaller internal set. There's a lot of good work being done and having been done already in the industry about representative scenarios. It's not at all unreasonable to expect a decrease of 80 to 95 percent in the number of scenarios you have to run for your base set. All those things could help. You also may be able to approximate the RBC piece over a smaller subset of model points than what you're using for your pricing.

You want to think in terms of having a compact model. If you're doing a living benefit, you may be used to pricing over ages 40, 50, 60, 70, 75, 80 or 85 because you want to capture all your mortality costs. But if you're doing a living benefit, maybe for that run they're not so important. Maybe you can get by with three issue ages because you're not pricing something that is directly impacted by mortality. It's only a second-order effect. Anything you can think of that helps reduce the length of your calculations would help.

We want to at least consider hedging analysis because, depending on what you're actually doing, you may want to use different types of scenario sets for that, too. There are going to be the market-consistent sets, which Dave is going to tell you about, that you use to actually put a market-based cost on hedging. However, once you've built a hedging strategy and you want to put it through the wringer, you're going to want to test it over a scenario set that reflects realistic possible outcomes so that you can at least analyze what happens to your tails. Do I reduce the tails that I have to deal with? Is my hedging program helping my capital costs go down (which is what you're partially after)?

With C3 Phase I you use fixed annuity book value and minimum guarantees, the guarantee that someone can surrender at book value or that you won't have a credit less than 3 percent. These are things that are going to involve a tail. It could be a double tail. The worst that can happen could be harder to define than VA guarantees, so you might want to analyze several kinds of yield-curve moves. But all this can help to put together a set of real-world scenarios that can help you analyze what your tail is. If you're in the UL world, again, you want to think of what your guarantees are, whether it's secondary guarantees or interest rates that you're most worried about.

The new GAAP Standard of Practice (SOP) 03-1 insists on using scenarios to cost your guarantees, at least a meaningful, deterministic set, if not a stochastic set. It's not imminently clear whether the costs should come from the mean or the tail, but it is at least an application you might think about in terms of what your tail costs are, and if it is in the tail, possibly not deep in, anyway.

Does your reinsurance curtail your tail? We're going to talk about hedging. You could have reinsurance that you think is going to help a lot with C3 Phase II; a lot of reinsurance plans cover insurance risks beyond a given point. I'll list a few examples. However, many lately have a coverage that has an upper bound like $10 million in a given year, or something like that. If the coverage has an upper bound and then you have to cover everything beyond that, are you really covering your tails? You might find in your C3 Phase II that you would be better off with a hedging program than that reinsurance plan because most of the CTE 90 is out beyond where their insurance stops.

In summary, one's definition of "real world" can be very much impacted by your relevant viewpoint on the "world." Real-world scenarios have many advantages in pricing for and discussing product benefit costs, particularly in the tails. Pricing guaranteed riders for VAs, but also pricing other products, can involve combinations of two or more sets of real-world scenarios. Pricing of those benefits should include both expected benefit costs and capital costs. These are both largely driven by tail results.

MR. DAVID M. WALCZAK: Let's start with a quick outline of risk neutrality and how we look at the tail risk in a risk-neutral context. I'll start by defining and contrasting risk neutrality with a real-world paradigm. Then I'll go to a case study, which we'll call "The Actuaries Go to Wall Street," and guess who gets popped in the nose on this one? Then I'll talk about a concept that's relatively new in the United States, but it's big in the United Kingdom and Europe. It's called "deflator technology" to value tail risk. Finally, I'll give you a couple of pages of references, stressing the fact that in this short of a time period, it's difficult explaining a topic that isn't really popular or popularly used. This is just going to be a taste of where you can go to find some of these models and references to risk neutrality.

Has anyone written a risk-neutral model, or does anyone use a risk-neutral model at work? There are a couple of people who are familiar with it. Let's talk about the fact that risk neutrality, in its purest form, means in that particular case the investor has a risk-neutral utility curve, or believes that every asset in the asset class returns the same. What further defines risk neutrality? Let's say that all investments produce the same return. A risk-neutral investor has a constant, or a linearly sloped, marginal utility curve. The options market makes use of this principle, and it insists that a dollar of cash and a dollar of bonds are worth the same at time 0 and that they should return the same for the options market to calibrate properly. Finally, we've got a derivative price distribution that's called risk neutral if a security price is defined by the return of the risk-neutral rate or risk-free rate. You can decompose the distribution into several states with observed returns "u" and "d" that solve this risk-neutral fundamental equation, and that's called a risk-neutral derivative price distribution. There's a lot more in the financial economics textbooks on that.

Let's continue with a real simple case. We're going to start with an example used in the financial courses from the SOA, the binomial lattice approach. If a Treasury bond that has no risk at all is guaranteed to return $5 on a $100 initial investment, then what do we do with an equity that could end up in two possible states, $95 or $125?

Let's say that we use the actuary view (or the real-world view) where we would build this out with 1,000 scenarios and equally weight them and start talking about the mean as a realistic way to estimate value. In this case, the actuary view gives you an answer of $110 for the equity, and there's a risk premium over the risk-free Treasury case on the bond of 5 percent since the equity is returning 10 percent and the Treasury is returning only 5 percent.

What happens if we price a call option based on that equity and that risk-free rate of 5 percent? We've got a 50-50 chance at a strike price of $100 of ending up with $25 of in-the-moneyness, or zero dollars of in-the-moneyness on a 50-50 propositional basis. The price of the call option is 25/2, which is the 50 percent chance of getting in the money, times $1/(1+i)$. In the actuarial, or real-world context, it's not always clear what you use for i. In other words, what's the proper discount rate?

Here's what Wall Street says is the right answer. Instead of trying to guess at an i value, you force the equity to have a probability distribution that's not actuarial or evenly weighted. You give the cases the weight that solves for an equity return of 5 percent. In our case, it's one-third and two-thirds. The call option price is different. The 25 is divided by 3 instead of 2 because there's less of an in-the-money chance in the Wall Street paradigm. You've got an interest rate of 5 percent dictated by the risk-free rate, and so there's no guesswork.

Let's look at some key implications of risk neutrality. We know, again, that all assets return the risk-free rate. That's how we were able to solve the price of the call option to match what the Wall Street tradable value would be. The subjectivity is gone from the discounting decisions, so we no longer have to guess at an appropriate discounting rate, which is the actuarial conundrum. You can project future risk-neutral rates. It's difficult. That much is clear. You match future prices, but not necessarily future return expectations. As Doug Robbins repeatedly pointed out, you're not going to meet a chief investment officer anywhere with risk-neutral expectations. I met one once who said, "I don't care if we're in Treasuries, risky corporate bonds, equities or pork bellies. They should all yield the same, but we are subject to the NAIC basket clause." There are future return expectations that every investment manager has.

Risk neutrality is arguably good for pricing and valuation of liabilities or assets that have tradable counterparts on options markets or equity markets. Risk-neutral scenarios as the paradigm is not necessarily good for projecting economic outcomes for earnings and EV volatility, nor for deferred acquisition cost (DAC)/value of business acquired (VOBA) amortization where stochastic processes come into play. There are a lot of reasons for that, which go beyond the scope of what we're talking about today.

Let's say you're one of the many multi-nationals out there that has francs or pounds in your world-wide corporate model, in addition to dollars. When you add a multi-currency parameter to the mix, the risk-neutral paradigm is going to break down because you're going to have currency exchange rates and different risk-free rates for different pockets of the company. Also, you can grossly reduce run time by using a risk-neutral paradigm. Doug talked about the stochastic-on-stochastic issue. With risk neutral, you can choose a small subset, run quickly and match Wall Street portfolio prices. That makes it a very attractive option.

Finally, when we talk about the Hardy regime-switching lognormal (RSLN) distribution, some interesting things come into play when you talk about risk neutrality. If you look at the Black-Scholes method, there's no return parameter that drives the cost of an option. It's simply the risk-free rate and volatility as the only two moving economic parts.

Let's look at that in RSLN paradigm. In RSLN, you're given two parameters that define the return of each regime. They are the good and the bad regime and a matrix for moving between the two. What if you have 10 percent, 25 percent as regime one and 15 percent, 25 percent as regime two? Your return percentages are different but your volatility curves match. These regimes are going to be equal in an option-pricing context that's risk neutral. If you're using that to price a call option, it doesn't matter that there are two regimes.

Deloitte has a Web site called The Smith Model (TSM)that's U.K.-based. On it is a public domain risk-neutral package. It's Excel-based with a VBA add-in that you can't get at. You can't get it and decompose the formulas, but it's a public domain and free. It will produce risk-neutral scenarios for you and deflators, which we'll define momentarily. It's great because the VBA function is kind of a poor man's ALFA or MoSes or Prophet, and the Excel is used for the reporting of the scenario results. There are Black-Scholes functions in MoSes, ALFA and Prophet. It's nearly impossible to easily program a risk-neutral (RN) equity scenario generator, but Cox-Ingersoll-Ross has an easily programmable one-factor interest rate model that's RN.

Let's continue with the practical uses for RN pricing versus real world. You can replicate actual tradable option prices in the option markets. Your reinsurer or financial counterparty in a hedge transaction is going to be pricing this way, so if you're not, you're going to come up with a different value and you may be taken advantage of or not understand the outcome. Embedded value is an interesting case where we might be moving toward risk neutrality because international accounting standards are pushing European multinationals toward a probable 2008 or 2009 use of fair value or market-calibrated value for embedded value. Some of the multinationals are already scratching their heads on how they're going to do it and how they're going to explain that potential massive increase in their liability cost.

The variable annuity example is simple. It's a single premium with a stock/bond fund allocation. The guarantee grows at a constant value every year. There's an income benefit (IB) and a death benefit (DB). There's one decrement of 10 percent a year. The policyholder receives the maximum value for a DB and survivors receive a maximum of a current base pay-out factor and the guaranteed basis.

We used 2,000 scenarios. We have a 15-year maturity period, and there's a 20 percent accrued guarantee. Right now the fund value is lower than the guaranteed value by a factor of 1,200, and we're rolling up at 2 percent. We need to find the GMIB pay-out rate at 10 years and 4 percent guarantee, and we've got a stock/bond mix of 60/40 in this case.

What did the actuaries and the Wall Street folks come up with? We saw a hockey stick-like distribution for the DB. As we go along the percentiles and into the fat tail, we saw that that's the distribution of the scenario results. The actuarial answer would say about 60 basis points is our realistic cost of the benefit. Wall Street, using The Smith Model public domain software to value the benefit, came up with an answer of 90 basis points.

Has anybody tried to obtain reinsurance for a GMxB in the late 1990s? (One person in the audience responded.) Did you perceive the cost of the reinsurance that was presented to you as onerously expensive in that case? The reason is that the reinsurers, early on, at least several of them that priced it properly, used a Wall Street counterparty-type approach where they dynamically hedged and came up with a risk-neutral answer that was much higher than the actuarial mean. You had this disjointedness in VA product actuaries and the reinsurance actuaries that were trying to come together and strike a reinsurance arrangement. Most weren't able to do it.

Let's talk about the GMIB next. There is much more of a hockey-stick pattern based on this benefit design. There is an extremely fat tail. There isn't any cost at all until you get up to about 87 percent and, again, the actuarial answer for the mean and the Wall Street answer differ widely.

The risk-neutral cost of equity optionality differs from the real-world cost. That should be your number one retention point from this talk. Bring away from here that there are different paradigms, and use the risk-neutral cost when you're dealing with a situation where you're being measured against an option market, a tradable equity market, or some kind of fair-value and market-value paradigm. You should be using the RN paradigm. In today's environment, because RN is a direct function of forward implied volatility, the RN equity cost on optionality is higher than the real-world cost. Maybe that will change as the equity markets flatten out indefinitely and the tradable cost of forward implied volatility drops.

Another explanation is that the RN method more heavily weights higher-cost scenarios. Actuaries dealing with Wall Street hedging and reinsurance types were getting the wrong answers and charging too little for the first generation of GMxB fees. Wall Street types believe that actuaries were underpricing GMxB fees because they weren't risk neutral. Models are very sensitive to implied volatility and policyholder behavior. If you're in a risk-neutral paradigm, a robust policyholder-behavior function along with the proper forward implied volatility for your case are the two crucial driving points of your model.

Let's introduce deflators as a risk-adjusted discount factor. It's a path-dependent, scenario-dependent, time-dependent discount factor. Let's say you're in stochastic scenario seven, year 17. There's a unique deflator out at that point that when applied to a cash flow, gives you a discounted risk-adjusted present value. Deflators could be designed to be real-world or risk-neutral. I'll explain an example of each one and a place where you can go to get public domain generated deflators to read more about it. The effect is to match risk-neutrally valued option pricing to a scenario set basis. Alternatively, you could come up with just the risk-free rate and a set of risk-neutral probabilities. That's all you would need because if you got $1/(1+i)$ to the "T" using the spot yield curve and a set of risk-neutral probabilities, then you can derive a present value of a risk-neutral distribution using the mean.

The definition of the deflator is that it's a state price, which is an ultimate price for cash flow, divided by probability. The conditions are that deflator at time 0 equals 1, and the future value has to be non-zero. There's a cap on the deflator, a capital asset pricing model (CAPM) deflator, from the capital asset pricing model that you can derive for a real-world or risk-adjusted deflator. You know from risk neutrality that you have got to get rid of everything past the risk-free rate (R_f). Well, the CAPM says there's a lift off from R_f. You can rearrange terms to try and solve for a portfolio return and then you can isolate X_0, or the starting state value, and solve for what's called a "CAPM deflator."

There's a lot more on this in a paper that is going to be required on the SOA Course AB starting this year for the first time. There are three Deloitte colleagues in London that have written a paper called "Modern Valuation Techniques." It introduces real world versus risk neutral and the use of deflators, both risk-neutral and real-world.

There are some good places you can go to find out more about risk neutrality. They include TSM and the Deloitte Web site, which has free stuff on it that you can use in Excel at *www.thesmithmodel.com*. Tillinghast-Towers Perrin has a CAPLink site, and there's an RN version of CAPLink that you can look at. There's also an american.edu Web site that has what's called the "Risk Neutral Density Function Estimation Toolkit." This is a way of coupling the spot curve with a set of risk-neutral probabilities that you can develop with a toolkit and come up with risk-neutral scenario valuation. Finally, there are some references to RN literature. As I mentioned, there's the "Modern Valuation Techniques" paper from the Deloitte group. The Hull book, "Options, Futures and Other Derivatives," talks a lot about RN versus real world. The SOA Investment Section's newsletter has a paper written by a Deloitte colleague called "Deflators: The Solution to A Stochastic Conundrum." Finally, there's the CIR model, which you can program in Excel. It's a one-factor 90-day model that the yield curve is fit through, and it's a risk-neutral arbitrage-free generator.