RECORD, Volume 30, No. 2*

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

Session 31PD
Hedging Variable Annuity Guarantees: A Practical Discussion

Track: Risk Management, Investment

Moderator: FRANCIS SABATINI

Panelists: MARK EVANS
DANIEL D. HEYER
DAVID HOPEWELL
ZAFAR RASHID

Summary: A number of insurance companies have recently implemented or are in the process of implementing variable annuity (VA) hedging programs for guaranteed death and living benefits. The panelists discuss the benefits of hedging these guarantees, the challenges they face and how they keep score.

MR. FRANCIS SABATINI: I have the pleasure of being your moderator today, so I have the distinct pleasure of introducing the panel. I think there are some common themes among the panelists this year. All four of them at one point in their lives were actuaries.

I'll introduce them in the order that they're going to be speaking. We'll start off with Zafar Rashid. He's the executive vice president and CFO of the annuity business at Prudential, and in the middle of implementing a hedge program. He'll provide you with some perspective on the business and the issues one has to confront when implementing a program.

Speaking second and more generally about hedging VA guarantees is Mark Evans. He's vice president and actuary at AEGON USA, and basically leads the equity-based derivatives program at AEGON.

He'll be followed by Dan Heyer, who is a property and casualty (P&C) actuary at Nationwide Financial. He runs Nationwide's guaranteed minimum death benefit (GMDB) hedge program. And then finally, there's David Hopewell. Like Mark and Dan, he's a senior manager at Ernst & Young, formerly a derivative trader at AEGON USA. Mark replaced David when David left, and David is now with Ernst & Young. All four of these panelists in one way or another are actively involved in trading derivatives. I think it should be a very interesting discussion. And with that I'll turn it over to Zaf, and he'll take it from there.

MR. ZAFAR RASHID: I have hedging VA guarantees. It has to be the most exciting topic in this meeting. When you're thinking about hedging some of these benefits, it's unplowed territory, where we're exploring a lot, looking for creative ways of solving the problems. The search for the magic bullet never ends, but there are probably few magic bullets out there.

Let me start with a little bit of background on Prudential. It's an old company, as you all know, founded in 1875, and has been a public company for about three years. The asset base is $454 billion with $24 billion in market capitalization. Last year, Prudential bought the American Skandia Company and combined American Skandia business with Prudential annuity business. That's when I joined the company, as CFO for the combined operation. That was just a great acquisition from Prudential's standpoint. I think the timing was perfect. It was right at the bottom of the market, and we've had great results since. The total combined operation now is close to $48 billion of annuity assets under management. These assets are primarily VAs. We're selling through four distribution channels: independent planners, Prudential agents, wire houses and banks.

For the benefit of those who may not be intimately involved in these things, we have just about every variety of GMDB. Generally, they tend to fall into four classifications: return of premium; step-up or highest anniversary value; roll-up that increases on a regular schedule of a few percent a year; and one that gives the customers the higher of step-up or roll-up. As you can imagine, the price tend to go up as you go down that list.

In terms of guaranteed minimum withdrawal benefit (GMWB), it's probably the newest of the living benefits. It offers the customers the guarantee that as long as they take their money out in small increments, they will at least be able to get their principal during a period of time. The most difficult is the 7 percent withdrawal guarantee. They can take 7 percent out every year, which means that if the market stays level during a 14.25-year time frame, basically, they'll be able to get all their principal out. Of course, if markets drop, then they can start receiving benefits, because they will have run out of account value and then it's up to the company to cover the rest.

The guaranteed minimum income benefit (GMIB) is somewhat similar, except that you have to annuitize to take the benefit. There are similar types of roll-up features. GMIBs stand to have somewhat higher amount of interest rate exposure to them because their interest rates are embedded in annuity purchase rates. You have to worry about that as well.

When we think about the VA guarantees, one can step back and look at what's happened during the last several years. Early death benefits were either free or offered at a nominal cost. I think the mentality was markets will continue to go up, and as markets continue to go up, you're never going to have to pay a claim, so the risk is minimal. Actuaries typically tended to price on that kind of a basis. The events of 2001 and 2002 were a big wake-up call. We all realized markets couldn't continue to go up forever, and that has had a huge impact. It's caused the reinsurance market to dry up, because reinsurers supported a lot of these benefits.

Once the reinsurance market started to dry up, the direct writers had to look at other alternatives. Not finding any capacity, they had to start increasing some prices. But it ultimately led to everybody wanting to turn to the capital markets. The capital markets are always there. They're not going to be here tomorrow and go when the capacity goes in and out of the reinsurance markets. So that's where everybody turns to, and of course, capital markets pricing is very different from the way actuaries have historically done pricing.

If you look at the difference, typically actuaries tend to look at historical experience both for the customer behavior as well as market performance and base pricing on that type of a model using cash flows discounted at a risk-based hurdle rate. What typically happens with these kinds of assumptions is that you come up with lower valuations of the embedded options buried in the product itself. Capital markets, on the other hand, are based on principles of no arbitrage, efficient behavior and a risk adjustment. There's a world of difference between the typical actuarial method and the capital markets method of pricing.

So what's happening now is that with companies wanting to look at the cost of hedging, and building that into the pricing, we have to bring in some of the principles of capital markets into our regular processes.
One of the things when you're trying to create a hedge program, if you're a public company, you have to be very concerned about GAAP earnings. We're going to talk about GAAP earnings and accounting treatment on different products, as well as the capital and other issues that you have to deal with when you're trying to construct a hedging program around these benefits.

Let's start first with the GMWB. It's very straightforward. FAS-133 applies. You have to take the embedded options in your VA product and mark to

market in every single accounting period. So the change in the value of the liability has to flow through your net income statements. What happens is that the volatility in that movement, because that valuation has to be done on capital markets assumptions, falls straight through to the bottom line if the benefit is unhedged. Now what you want to try to do is create an offsetting hedge that will have ideally an equal and opposite movement. The benefit of that is that the liability and hedge asset are marked to market through the realized gain and loss line, which is a single line. It's usually not part of operating income, unless you define it as such, but is part of net income. So the fact that it gets reported in single line, and assets and liabilities are marked to market at the same time, makes it somewhat cleaner from a presentation standpoint. That's deceptive. As I can assure you, it's not that straightforward. The methodology for how you mark to market is still in the state of development.

The state-of-the-art is not completely settled. I'll mention that we're working with three major accounting firms on these issues. I can assure you that there are great differences of opinion between the three, and there are even differences of opinion within our primary auditor in terms of how you should go about valuing some of these methods. So this is an area that is changing quite a bit. I'm sure there will be more emerging in it, but at the end of the day, from a risk-management standpoint, it's probably more straightforward than the GMDB and GMIB.

Let's look next at the GMDB and GMIB accounting treatment. These benefits are specifically excluded from treatment under FAS-133 because they have embedded in them mortality characteristics that are deemed to keep them out of treatment as an embedded derivative. So what that means is that you don't have to mark to market, but SOP 3-01 applies to these benefits.

Now, the SOP 3-01 is also in the stage of development. It's not completely stable, and there's lot of discussion going on in terms of what kinds of stochastic models you need and how you go about adjusting those models for updating and computing your reserves on an ongoing basis. But the net effect of it is that it's essentially an accumulation of retrospective net fees for these benefits. They determine a benefit ratio and that percentage of your fees are accumulated. Now for most of the benefits, your claims are going to be way out there, so when you introduce the benefit initially, all you have is basically an accumulation of fees. With claims coming out against it and as the business progresses, then it becomes the typical retrospective methodology. You do have to unlock your fees when assumptions change, and there are different ways of offsetting your unlocking methodology and parameters and corridors and so forth that you use in setting your unlocking approach.

One of the other things you have to worry about is the capital impact. I'm

sure most of you have heard about the Phase 2 for risk-based capital (RBC). Essentially, if it's the direction that the NAIC is going, the capital required for these guarantees would be essentially a modified CTE 90 amount of capital. So, what you need to determine is the CTE 90 quantity, subtract from it the reserves that you're holding and that's the amount of required capital needed for these products. That will be an improvement from where we are now, where we're using guidelines and some static factors that have been developed. But there's a lot of work that needs to be done to quantify what that impact is going to be and how that will affect the kind of capital requirements that are generated by some of these guarantees.

We've done some testing with different types of both these benefits. You can design your product to somewhat mitigate the end exposure. The real impact of RBC Phase 2 comes in when markets move up and down, and as your benefit goes in the money from being out of the money. It can be a substantial increase in the amount of required capital that you need to set up for these benefits. And that's something that in itself would create enough pressure for companies to either reinsure or hedge the risk, so that they're not exposed to that kind of volatility in the amount of required capital that they're going to need under Phase 2.

Having said all of that, it sounds pretty straightforward, but I think a lot more work still needs to go on in terms of some of the details of how that capital is going to be computed. More importantly, if you do have a hedging program for guarantees, how much credit would you be allowed to take into your capital requirements for having the hedge? In other words, can you completely offset the movements in your liability with movements in your assets? Is there going to be a minimum floor below which your required capital will not be allowed to drop? Those are the kinds of issues that still have to be addressed. That's in a sense what makes this challenging because while all these things were going on, we still have to create the products, market them, price them and take on these risks. These are risks where we're locked into them for the life of the product. It's not something we can offer and say, we didn't like the way we did it and withdraw the risk on existing business three years or five years down the road. So that is one of the challenges.

Now from a capital standpoint, the volatility of this required equity under RBC Phase 2 wasn't enough to create problems. You also have the issue of statutory accounting. If you have derivatives in your asset book, under statutory accounting, they get mark to market treatment. If nothing else is done, if you're holding a portfolio of options on your statutory balance sheet, each quarter, you would be marking those options to market. If there isn't a corresponding movement on the liability side, that volatility in itself flows through to the bottom line. To address that, you have to look at Statutory Accounting Principle No. 86. That defines the criteria that must be met to get hedge accounting treatment for any hedge used for these benefits.

If you can get hedge accounting treatment, then you don't have to mark to market the derivative, the options that you're holding every accounting period. You can amortize them in a manner consistent with the movement of the underlying liability, the liability that is being hedged. That sounds very nice, but doing it is not the easiest thing in the world. Rules around what you have to do in order to qualify for hedge effectiveness are borrowed out of the GAAP rulebook. With the GAAP rulebook, I don't think people realize quite what it meant to make it operational and actually do it on a day-in and day-out basis. You have to be able to designate every instrument that you buy as a hedge for a specific liability and you have to be able to prove up front that the hedge is effective within a certain period.

Then you have to be able to prove after the fact, that yes, the hedge was with hindsight, effective within a certain *period*. There are a whole slew of rules that you have to follow to keep that hedge effectiveness treatment and to get the stability in your capital position and statutory accounting statement.

And if that wasn't complicated enough, you really can't get statutory accounting treatment for VAs at the direct company because of how much is involved from the GAAP world. You essentially have to wrap it in a reinsurance contract and use reinsurance as part of the mix to make it work. I don't want to sound like I'm making a pitch for the reinsurers here, but I think that's what I'm probably doing.

So with that as background, I tried to list some of the thought processes that we had to go through in thinking about creating a hedging strategy. The first thing you have to do is define your hedge objective. What is it that you're trying to accomplish? Are you trying to neutralize the economic impact? Are you concerned about capital? Are you concerned about GAAP earnings? Because the type of hedge that you structure will vary significantly depending on the objective. At the same time, you need to define what guarantees you're going to try to hedge and establish a risk tolerance.

Establishing a risk tolerance sounds very easy. Of course, you ought to be thinking about it and doing it, but how many people have tried to ask their management what their risk tolerance is? Let's have a show of hands. How many people have gotten the straight answer? I don't see a single hand. It sounds like something we as insurance companies should have been doing for 100 years, but I can assure you that it is not that straightforward. Somehow you have to grapple with the issue of risk tolerance. What types of risks and what magnitude of risk are you willing to tolerate? The automatic answer I get from our management is, "Well, we don't want to hedge every risk. We have a large balance sheet. We can tolerate a certain amount of risk." And that sounds right, but then the question is, "Okay, now how

much?" That's when it becomes a more complicated issue.

Let's assume I've defined my hedge objective, and then the guarantees to hedge, and established a risk tolerance. Now I have to create hedge strategies and slap a strategy together out of the multitude of choices. This is where it becomes complicated because there are so many choices and it's so hard to choose between them. You have to evaluate the effectiveness of the hedge in relation to the cost of the hedge.

There are some operational issues you have to deal with. Are you going to do dynamic hedging or static hedging? The breadth of hedge instruments you're going to use is another operational issue. Developing appropriate models and tools for valuing these hedges is a pretty complicated exercise. When all that is said and done, you have to establish the management procedures for managing the hedge on an ongoing basis. Who is going to be involved? Who's going to make the decisions on what to buy and what not to buy, what to hedge and how much has the liability derivative moved, what compensating trades do you need to make? On top of it all, you have to document all those procedures for your own internal sanity and for clean opinions from your external auditors. There are a whole bunch of operational issues that have to be dealt with.

Finally, when creating a hedging program, it's really important to consider your overall equity exposure. It would be inadvisable to be looking at only the equity risks and only in one or a handful of products or product lines. You really have to step back and look at the whole organization and say, "Where are all of my equity risks? Am I satisfied with the mix across the whole organization?" At the end of the day, that's really what is going to affect shareholder value, not necessarily the individual components.

You have to consider the overall interest rate exposure, and embedded in all these liability derivatives is a significant amount of interest rate risk. How does that interest rate lie with your asset-liability matching (ALM) exercise of the other interest rate exposures that are sprinkled throughout your organization? I mentioned risk appetite already—that we have to grapple with the issue of: How much risk are you willing to take? Implementing a full and complete hedge that completely neutralizes all the movements is probably not the right answer. Nor is a completely dynamic hedging strategy that is only hedging the immediate marginal impact of market movements. The right answer probably is somewhere in between, but I have to show you from some personal experience that it's a lot easier said than done. It's easy for me to come up and talk about it, but when you actually try to do it, it's a great deal more complicated. I'll leave it up to Mark and the others to show us how we're going to do that.

MR. MARK EVANS: At AEGON I'm in charge of all the hedge programs that

relate to VA minimum guarantees, equity-indexed annuities and equity-indexed universal life. I'm going to jump right in and discuss some of the practical issues. A lot of these are, in a way, peripheral issues, but they end up being pretty important things that a lot of times don't get talked about.

One thing we pay a lot of attention to in hedging the VA guarantees is rho. In other words, the sensitivity of the guarantees is such that as the interest rate goes up, your liability goes down. And vice versa is true. Falling interest rates ensure liability goes up. We actually think that for VA minimum guarantees this is a bigger deal than vega, which you'll probably hear about more. So, in our hedging software, every time we run the model, we update it with the current swap rates. We use those swap rates to do our discounting and also for the expected growth in our equity portion.

We run the model every night. If interest rates change dramatically during the day, we'll rerun mid-day. We won't do as many scenarios but we feel we can get delta accurately in about an hour of computer time. And we have some tricks that allow us to do that. That I'll get to in a bit.
We also have the ability to use stochastic interest rates in our hedging model. We don't reflect that when we actually calculate the delta that I use to make my derivative trades with, but we do have that capability. The other thing we're careful to do is when we're calculating Rho on our fixed accounts, those all have certain durations and interest earnings are locked in for some period of time. When we go ahead and shock our interest rate curve, in order to support the calculation of rho, we're really careful that for the maturity period of the assets, we've already gotten our fixed accounts; we don't increase or decrease the earnings on those.

Something else that hasn't gotten a lot of discussion in the VA hedging world, but we've found to be particularly important, is interest rate convexity—the second derivative of option price with respect to interest rates. So we will use the stochastic interest rate capability, not directly to hedge with, but if we want to know what the cost of an embedded option is in one of our products. We'll use our hedging model, except turn on the stochastic interest rate capability and that way we can figure out what the additional cost of providing a guarantee is, from the convexity and interest rates. That becomes particularly important if you have, say, a living benefit, where you expect a lot of customer behavior. That just takes the convex of the curve and increases that curvature; sometimes dramatically.

There's some discussion of hedging to accounting. We hedge to, I guess you could call it risk-neutral cash flows, and then hope the accounting works out. But if you are going to hedge to accounting, there are some things to watch out for. Make sure that your hedges will provide a reasonable protection from the actual cash flows that will result in a bad scenario. Another thing you have to spend time worrying about is that you're trying to hedge your GAAP

accounting or your statutory accounting. With AEGON of course, we have international accounting to worry about since we're not a U.S. company. These rules are changing all the time. Something else you want to watch out for under classical risk-neutral hedging is that the things that you're hedging to tend to move in a reasonably smooth pattern. In other words, as your stock market moves up and down, as the interest rates move up and down, your hedging under a risk-neutral hedging gains or loses in a reasonably smooth fashion.

If you're hedging to accounting, you want to be careful that that approach keeps that smoothness or at least that it's not distorted too much. I've seen some proposals on ways to hedge the accounting results that have you in a situation where you would have to suddenly change your hedge position dramatically for just a small movement in the market. And that's likely to get into some extremely high hedging costs if you happen to have a half-year that might otherwise seem benign but causes you to do a lot of rebalancing.

Here are some pesky modeling details. First, I want to mention rebalancing. I'm not talking about the rebalancing that goes on with dynamic hedging. I'm talking about customers who sign up for automatic rebalancing programs. It does funny things to your hedging results. So it's important to try to reflect that in your model.

Systematic withdrawal behavior is important. You should try to capture that from inforce data and identify those customers who are on systematic withdrawal and reflect that. For qualified business, you want to concern yourself with minimum required distributions; model that. Next are dollar-cost-averaging (DVA) transfers. We get into our hedging software. We capture the DCA amount and estimate where that's going to end up and project those transfers in our models because if something is not an equity now, but it's going to be at some future point and you know about when that is, that does affect your hedging.

There's been some talk about static versus dynamic hedging. I just want to make a point from a practical point from derivatives trading. If you're trying to hedge with only puts, you're likely to run into trouble. Because sooner or later, you're going to have some change in the asset or the liability, something you didn't expect, where you're going to need to sell that put. Especially long-dated puts are just frankly not that liquid. So if you're looking at hedging with puts, that's fine. If you want to go into a static approach, you're probably better off having some combination of puts and futures. Futures are liquid, very easy to trade, and you want to take that into account when you're setting up your hedging program.

Attribution is very important. Set up an attribution system so you can compare your hedging results to the movement of your liability. You have

your capital market variables that you want to track. For example, you want to make sure that the cash flows and the change in the value of your derivative instruments is matching up with what's going up on the liability side. It's also important to track your modeling assumptions, mortality, lapse, and whether customers are doing a lot of transfers back and forth between sub-accounts.

Need for speed: Hedging is a very computationally intensive process. If you haven't heard, I'll just tell you that there are some companies hedging blocks of business with dozens of PCs running every night, all night long. We've come out with some tricks that help us get around that. We have very large blocks that we did with two PCs, and like I told you earlier, if we have to, we can get accurate measurements of delta in about an hour.

One thing I want you to take away from this presentation is that for hedging purposes grouping is bad. It's bad and it's unnecessary. A lot of the relationships are non-linear. In a lot of the groupings that actuaries have typically used, things just kind of average out. In the option, because of optionality and because of the way strikes work, etc., you don't get a lot of relationships that are linear. That makes it hard to group. If you do go through a grouping exercise, then you have to validate that your grouping hasn't distorted something and be extremely difficult. It's not only difficult; it's a waste of time.

The reason that it's a waste of time is, for hedging purposes, what you really want to do is use a different set of random scenarios for every single policy. For a lot of actuarial-type calculations, CTE calculation, you have to, for various reasons, use a common scenario set across all policies or all groups. For hedging purposes, each policy has its own scenario. Add up the results from all your policies and what you'll find is for a lot less computer time, a lot purer total scenarios, you'll get much more accurate calculation to set your hedge positions.

Another trick we do to speed up convergence and run time is on larger policies. For really small policies, we may actually only do one scenario. There are some other things we'll hit pretty quickly. You might as well use antithetic. It's pretty benign; it's pretty safe; it'll give you some extra speed. The other thing is because everything is offset, it allows you to check your distributions more quickly, to make sure you don't have any little bugs in your scenario projections.

Since we calculate different scenarios for every policy, we calculate a scenario, use it, project it, save the answer, throw the scenarios away and just save on computer memory. Otherwise, we'd be saving literally tens of millions of random numbers every night.

We do our calculations by shocking scenarios up or down by either the equity level or the calculating vega, and we have random customer behavior in our models. We're careful to make sure that the random numbers are the same for the shock scenarios as well as the base scenarios.

If you're building a hedging system, it really takes a combination of actuarial knowledge, the derivative knowledge and the software construction knowledge. The optimal thing to try to do is find an individual or individuals who know all three of these, and leverage off those individuals to get your system built. Of course, there is the age-old question: Do you build versus buy? It obviously depends on the resources you have available. If you don't have all three, you probably want to look at buying. Of course, if you build, you can do things to your specific products. I can say I'm just trying to talk about some of the practical issues we face and have to resolve to get our hedging programs in place.

MR. DANIEL D. HEYER: Before I start, I have three disclaimers. The first is, throughout the presentation, if you have a question, please ask it. I won't be smart at the end of the presentation and you'll save me the embarrassment of forgetting what I was talking about. The second is, I'm not a life actuary, but that doesn't stop me from using life actuarial terms in non-standard ways. I drive my coworkers crazy; please don't let me drive you crazy. Finally, there's going to be a lot of overlap because we're all carrying the same brick. There are a lot of nuances about how you can drop that brick on your foot. So parts of this I'll go through very quickly and highlight what's different without going to the basics of it again.

First of all, what is the GMDB? In effect with a GMDB, if you die, you're going to get at least some minimum amount of money back. What makes this good for the policyholder and really bad for the insurance company is now we're exposed to two sets of assumptions. We have both the actuarial assumptions of mortality and policyholder behavior, and we also have all the financial valuation assumptions about equity, volatility, interest rate environment and things like that. But fortunately for us the financial risk is freely traded, so that can be a block of risk that we can sell back to the market and manage how much we actually retain.

I'm not going to go into all the different types of GMDBs. Zafar did that very well. But I'm going to try to focus my talk more around valuation, how you value these things and think about them. To do that, we need to look at the different aspects you would want to value from a financial standpoint. The first thing you want to do is break benefits into two different types. One type we're going to call vanilla, and these are things that are essentially one-dimensional. They're like a put. You tell me what the equity index is at maturity, I tell you what the payoff is. I only need one-dimension to fully characterize the payoff.

When you start getting into fancier things like reset and ratchet, where the level to strike, the level of protection can actually move up if the market moves up. For those, all of a sudden we have added dimensionality. I have to worry about what the level of the equity index was at all these different book bases, before the maturity of the option.

Finally, a ticklish thing you always have to deal with is that these are basket options guaranteeing the aggregate performance of multiple equity and interest rate instruments, and not just the single index.

We talked a little bit about GMWBs. How is a GMDB different? In a nutshell, you have to die. So that takes away one piece of optionality for us. The policyholder really can't choose how he's going to exercise this thing. You also have a slight difference in what I call monetary lapse rates or attrition of account value. A GMDB is exposed to money staying around. A GMWB is exposed to money leaving.

Now we get into kind of the nuts and bolts of thinking about what it is that we want to value. In its purest form, what we want to value is the payoff diagram. We have probably seen that for a put, where you put the equity index on the axis and the payoff on the vertical. When you start to think about a GMDB and all the different dimensions in it, it gets too complicated to really think about a payoff diagram. A simplifying way of visualizing it, you don't take it too far, is to think of a GMDB as being a ladder of puts and then for each rung of this ladder, you're going to have: What is the strike? What level of protection am I guaranteeing for that maturity? And how many do I need to buy? The product design is actually controlling the put strike. If I have a roll-up versus a return of premium, obviously, there's a different level of protection.

Finally, the actuarial assumptions control the number of puts. How many people are going to die to actually be available to collect? Are those distinctions clear?

Now we get into two classes of real unknowns. Whenever you hear anybody talk publicly about hedging, they're always saying, "Yes, we include policyholder behavior in our hedge program," but it's really devilishly difficult thing to think about. Here we talk about investment behavior.

You can take the simplest case: The policyholder puts all his money in the S&P and walks away and leaves it. In that case, you have a simple index equity put. At the opposite extreme, you have somebody who has his magic day-trading strategy. You have written in the guarantee that he's not going to lose money on the day-trading strategy, and you provided a passport option. You can get out your textbook and look at the difference in the option

price, an index versus a passport option, and it's absolutely a mess. The policyholder behavior from an investment standpoint can really make a big difference, the guarantee is in writing and the type of hedging you want to buy.

There are companies now that will manage people's VAs for them and trade them according to different strategies. Is that something you would want to try to hedge?

The other thing you get into is pure lapsation risk. In an up market, people might want to leave or stay, depending upon what the benefit is, what their option is worth, and how old they are. So if you go back to thinking about this ladder of options, essentially what you're saying is every time the market changes, you're going to have to redefine the ladder of options you're buying, because it's going to be a different number of people hanging around for each rung of the ladder. How are you going to handle this?

If you really believe your dynamic lapse model, you could hedge that. You could say, "Well, I know I'm going to have to buy new options here and sell some there, and I want to have that risk hedged. I have the money to take care of it."

The other thing you say is, "Well, when it happens, I'll deal with it. I'll just buy some and resell some and accept it as a risky hedge re-balancing cost." Or the other thing we didn't say, "Time zero I'm going to lock in my assumptions, that ladder's stuck and done and I'm not going to worry about it any more and just accept the mismatch." Again, you have to talk about what my risk tolerance is and what kind of an effect this has from a return profile. Then you decide how you want to handle it.

That was a whirlwind tour of some of the major financial things you have to think about, putting a price on it so you can manage it. When you start asking what a GMDB is worth, you get into the issue of in what framework, actuarial or market valuation, and this was touched on briefly earlier. Actuarial valuation is based on the law of large numbers. On average I'll come out okay. But financial valuation isn't based on that at all. It's based on manufacturing costs. Spending an amount of money today, if I manage it properly, I'll come out whole every single time. Those are completely different frameworks and they yield completely different prices. Intuition should tell you that. On average this is okay every time; which one is going to be more expensive?

I found that you always have interesting discussions with senior management whenever the market valuation is much larger than the company-specific valuation for the same risk.

Buy versus build. This was touched on briefly. There are good reasons to buy and good reasons to build. There are good reasons to both buy and build and frankly, there are good reasons to neither buy nor build. I'm going to go through some of the high-level things that I think about on this question. You're going to see that I have an obvious bias and there will be a quiz to see if you can tell what we do.

The first thing you're going to have in a model is a statistical description of the market. What do I mean by that? When we talk about Black-Scholes, we say well, equity returns are lognormally distributed. There's no term structure to interest rates. There's one universal interest rate. That's the statistical description of the market, nothing more.

Regardless of how you're going to implement your hedge, the statistical description of your market has to be able to reprice the both your risk and your hedge instruments. If I have a model that says, "Oh well the price of this put is $5," but when I go look at the net, it's $6. I have a problem. I don't really have any idea now in any sort of consistent way what the value of this thing is that I'm trying to hedge, let alone what the cost of hedging will actually be.

The other thing you run into is, well, we would like to have a model that describes everything. The more complex you make things, the more difficult it is to implement that model. The thing to keep in mind, if you're actually going to purchase this system, is that this is still in. When you're evaluating a system that you might purchase, you really need to understand what's under the hood and understand what sorts of risk you are and are not evaluating. And with the Black-Scholes model, you're not evaluating interest rate risk at all. That would be something you'd want to know if you're going to buy a Black-Scholes model.

So fine, a statistical description of my universe doesn't help me calculate prices in another cell. There is an array of ways you can go in calculating a price. When you're lucky, there's a closed form solution. In the Black-Scholes formula, this is a case of that where you have a closed form formula that you can use to calculate the price for that universe.

Sometimes you're not so lucky, but you can get information about the distribution of returns at different maturities, in the form of a characteristic function. You give me these numerical tricks to get a price that way. Or you can get even harder and you can say, "Well, I can take everything over here from the stochastic land and turn it over here into this deterministic-partial-differential–equation land and I can get prices that way." Of course, we can always go to the standby of Monte Carlo simulation. We all think of that as being easy and intuitive and straightforward, but let me be the first to assure you that there are some very ticklish things in there, that are very easy to

get wrong and can lead to wrong valuations.

I think the key thing is again to fully understand what the statistical description of the universe is using it. Do you believe that equity returns are lognormally distributed and that there is no term structure to interest rates? I don't believe either of those things. And in that sense, Black-Scholes is perfectly wrong. Now it uses a market quoting convention because it's a nice shorthand way to summarize a lot of information about how it's wrong, and it's all loaded into that volatility number there. When I look at other things, like strikes and maturity, then I can get information about how the market believes that Black-Scholes is wrong. So even though it can still be useful, and we have a lot of residual models around that we know are wrong and aren't what we'd like them to be, we understand how they're wrong and but we'll still use them.

We have these four main approaches: closed form solutions, numerical inversion of asset characteristic function, numerical integration of pricing partial differential equation (PDE) and Monte Carlo simulation. A formidable technical tool kit is required to be able to move between these four approaches. There are more approaches, but the payoff is absolutely nil. Mark talked about the need for speed. You can get three orders of magnitude difference in the time it takes to value an option if you match the right technique for the right benefit type. So, it's expensive and difficult to find these people, but to find the right staff, it can really help you tailor the right approach to the right benefit, which is very important. Purchase a system that will value the widest array of things. It isn't necessarily going to be the most efficient.

You thought the other stuff was hard; calibration is absolutely the hardest part. I have to actually calibrate the model. I have to estimate parameters from market data. Since all models are wrong, but they're useful in different situations, you absolutely must be very careful about the instrument you choose to calibrate to. For example, if I calibrate to vanilla puts and then try to price-ratchet structures, I absolutely guarantee you you'll wind up with a nonsense set of prices. The reason being that the put only cares about the distribution of returns at a specific time, but the ratchet actually cares about how you got there as it meandered through different periods. This is a completely different set of relationships that you need to be able to calibrate to. You have to think about this very carefully, both from the standpoint of a liability and the instruments that you're going to buy. If you can't re-price your hedge instruments, you're dead.

You also wind up with a problem in that the method that is efficient for generating your option values quickly might be very lousy for calibration. You also have this problem with the most important parameters (such as market volatility), which aren't observable. You can't really go look up what the

current volatility index is for your specific model. Calibration is something that takes a lot of time and a lot of attention to do it right. How often are you going to do it? And from what I've seen, sometimes this isn't even included in purchased systems. That's not particularly helpful.

I'm going to be a whirlwind, talking about hedge structure and attribution, and this is more about raising questions for discussion and some things to think about. I don't have any answers for any of these per se, but definitely would be happy to discuss any of them, either now or at the end. We already talked about what constitutes risk you want to retain. Textbook hedging discusses economic hedging or cash-flow hedging, but insurance companies are very much worried about GAAP and stat accounting. Is that something you want to try to hedge?

Conceptually, do you want your hedge to buy back the GMDB that you sold to the policyholder, or are you going to try to protect the profitability of the overall policy? Do you want to hedge the performance of individual policies or the aggregate block? All these things get into different types of structures that you could entertain.

Do you have enough faith in your models to actively trade very complex derivative instruments? If you trade simple instruments, you can have a lot of management involved. You definitely can't let things go on their own for any length of time. But if you trade complex things that do a nice job of defusing your liability, now you're betting that your model did a good job of picking the right set of things to buy. And then finally, does your investment department have views on the relative attractiveness of different hedging?

I'll reiterate what Mark was saying, as to the plan on seriatim valuation at first; careless bucketing will create absolutely immense distortions. You can bucket, but it requires a very deep understanding of the scale and variant properties of option prices in the construction of your portfolio. You can equivalently think of time to maturity and volatility as being interchangeable, and there are adjustments that you can make to map everything into the same maturity, by playing with the volatility that you use to value each option. There are a lot of other relationships like that, that you can leverage to, and actually collapse things down into the number of runs that you need to do.

Tracking and attribution: This will take just as long or longer than everything else we've discussed up to this point. Do not underestimate the amount of work required to do this. In tracking and attribution, you might cover six basic areas: the exposure that you're actually trying to value, the assumptions that you use, actual versus expected income, the tracking error, changes in liability value and finally, earnings or capital volatility. Yes, there are the issues like trading executions, quirks of your hedge programs. For

example, if your data system still isn't updated as regularly as some of your other information feeds update, you'll get real disjointed. They have to be tracked and attributed.

You are not going to hedge all your model parameters. Some of them you'll say you're willing to take that risk. Finally, tracking errors can come from any deviations from modeled assumptions. For example, Black-Scholes said the world is lognormally distributed. Well, what if it's not lognormally distributed? What if it's *iffy* for part of it and normal after that? That's a difference in the universe and what you thought it would be and that is going to flow through. You need to detect that and track that.

Finally, I have this quote of Matthew Prior hanging in my office. "From ignorance our comfort flows. The only wretched are the wise." I'm on the short end twice, because I'm both ignorant and wretched.

MR. DAVID HOPEWELL: I'm going to talk about hedging guarantees, that is, examples of portfolios that do the job. There will be a little bit of philosophy but nothing quite so profound as at the end of the last session, which I heartily agree with.

I would say there are three broad options: static hedging, full dynamic hedging and partial hedging such as a delta and/or rho hedge. Static hedging is potentially inclusive of dynamic hedging and partial hedging. I hope that at the end of the program that will be clear. I believe that. And as you might guess, each program is different. It costs different. Sometimes the cost depends more on the market, sometimes less. They require different interest structure, different thinking, different approaches and different expectations to run.

This is what I thought should be last. This is what I'm calling a partial dynamic hedge program. Many of the people in the room are familiar with delta hedging. Delta hedging is offsetting the sensitivity of these liabilities to the equity markets, through in the case of VA guarantees, selling short equity exposure. You sell short equity exposure through the futures market. You get the equivalent of selling stock and receiving cash and investing cash in the money markets. Delta hedging is the momentary conversion of an equity position into a cash position.

Delta hedging works. It doesn't work fully. The degree to which it works varies over time, which is the main criticism of it. When you think about a VA guarantee, at least I would think at first, it's all equity risk. The markets go up; you're happy. The markets go down; you're not as happy as the issuer. Well, it turns out when you do the attribution that has been referred to by at least one of the other presenters, delta is only about one-third of the total risk of these benefits. So whether the market is at 100 percent or 90 percent

or 80 percent of its starting value today, it is only a partial cause of the cost to that point. And for about the last year, that's been one-third. Rate risk has been about one-fourth.

We have had a pretty interesting rally up in rates during the last year. I think it was June 13 of last year when they troughed, after having come down about a point and a half. We're about a year past that. They've come back up. It was a pretty big year.

Next I'll discuss option volatility changes that are associated with vega. Near the end of last year, they were relatively normal because the markets had stabilized. Long-term option volatility and implied volatility came down a little bit. In the early part of this year, they fell rather dramatically, such that Wall Street firms were feeling pain for their inventory of options. When the S&P stopped going up and started to go down a little bit, option volatility firmed, but didn't really rise back to where it had been late last year.

Short-term equity and rate movements are the second moments of the distribution. That is gamma and convexity. Guaranteed benefits had some sensitivity to short-term volatility in the equity markets and some sensitivity to short-term volatility in the interest rate markets, which you'll see if you mark them to market. For equity markets, the short-term volatility may, probably does, presage long-term movements. For interest rates it's maybe more arguable. With interest rates, the short-term movements have not tended to be predictive in the same way the short- or the long-term movements or the same way equity movements have been.

One thing I didn't hear mentioned today (I'm surprised and pleased that I had at least one unique thing to offer the group) is that the term structure of interest rates matters to these benefits. That is, they have significant key rate duration exposures. So hedging just parallel shifts probably is not sufficient at this very least; you want to know at what maturity the hedge ought to be.

And in FAS-133, for years and years, it seemed like FAS-133 was meant only to cause pain to people selling guarantees. But in fact, when you're hedging there are other more painful accounting standards. SOP 03-1 may be one of them. Given the pliability of some valuation standards that seemed to be written in black and white, I'm sure companies will find ways to make it work for them, but that remains to be seen.

This is an example of a partial hedge.

Partial Greeks Hedge

Liability Greeks	
Delta	(46,845,846)

Gamma	173
Vega	1,664,330
Rho	(448,082)
Convexity	152,899

Hedge Greeks	
Delta	(46,750,000)
Gamma	-
Vega	-
Rho	(448,082)
Convexity	(20,145)

Instrument	Quantity
S&P Futures	(170)
10 yr Swap	-54,777,806

This is a delta hedge. It also is a rho hedge. If you look on here, delta is a big number; that's the equivalent of $46,845,846 of stock. In order to offset that, I enter into a hedge with $46,750,000, because my futures are integral and there was an amount of $275,000 each assumed in this illustration. There is also a rho component, an interest rate component. That $448,082 is for a 10-basis point move. That means this liability moves $44,800 per basis point. In order to offset that, I needed to enter into an amount of swap that has an equal movement under that basis point move, and it turns out to be received fixed for $54,777,806.

One of the things I have observed among those who read books on Delta hedging, and even as I have done it, there is some consideration of transaction costs. Books on hedging give ways to analyze or assess what the optimal rebalancing strategy is, given certain transaction costs. For futures, transaction costs are low enough and falling, and they probably are negligible. I guess you wouldn't want to rebalance continuously throughout a day, or every tick, but the prices of electronic execution are getting pretty low.

If you believe your delta, you want to use that low transaction cost or rebalance more frequently. If after knowing the transaction costs are negligible, one still doesn't want to rebalance frequently, it implies maybe a lack of confidence in the Delta or more likely the assumptions underlying it.

I'm going to go to the fully dynamic hedge program. In this case, we're not just going to futures; we're going to trade options as well. In this case study, a GMWB, the proper hedge is a basket of puts. There are features of withdrawal benefits that might require calls, but that's not in here. For a fully dynamic hedge program, the more you match, the fewer the market-driven surprises are. Now that's worth thinking about. Because aren't the Greeks

delta, gamma, vega, rho? Yes. Those are the Greeks that we're used to talking about, but there are a number of others. For many of these benefits, I think as Daniel pointed out, there are correlation terms which are second order Greeks, between underlying. There are also higher order Greeks, and there's the interaction of some of the first order Greeks with time. And all of those ideally would be matched or matchable. In fact, they're not. They're not very matchable. That challenges, at least in my mind, how effective a Greek's matched hedge can be at the end of the day.

A basket will include puts and maybe calls. Ratcheting GMDBs probably require some calls. Resettable GMWBs probably require some calls. Roll-up death benefits probably do not. GMIBs, probably do not. Oh, there is one other thing—you can match the Greeks with more than one portfolio. That's one of the real difficulties, I think, of a pure Greek's matching framework. There's not a unique solution to the problem. And these portfolios, although they may have similar sensitivity at time zero, market equals current, have sometimes startlingly different behaviors as markets change and time changes.

Ernst & Young and I know how to calculate a portfolio of options to match a liability. In general, you can match a liability gamma and vega using two options at most. That may be necessary, but it's not always sufficient and in this example, I use three because I knew something about the liability that gave me more information than was contained in the time-zero Greeks.

This portfolio comes very close with the options to being a full hedge. So in fact, I'm only $400,000 worth of swap and four futures contracts for $275,000 each, off on Delta. That's a pretty good hedge just with options. But that wasn't the actual intent of constructing this hedge. It is to match all the Greeks with options. It was a side effect of the construction technique.

I am going to talk about this. The Greeks describe the hedge. The Greeks do not prescribe. They are not prescriptive. They say there's more than one solution for a basket of options and a set of Greeks. Greeks are not the only choice. The Greeks are the messengers of the risk-neutral world. The risk-neutral world is a very effective and theoretically attractive way to solve option-pricing problems.

Options also exist in the real world. Options exist in the world of markets and cash flows, and they have deterministic functions that can be used in the world of markets and cash flows to also construct portfolios to calculate values and so on. Now I'm going to go back to what should have been there, and that is the static hedge portfolio.

I have to give credit to Daniel. He came up with a description of this that I really should have and wish I had. He called it "hedge to the edge." "Hedge

to the edge" is absolutely the right description for this approach. And this is based really on a number of years I had as a derivatives trader. I finally had to give it up. I decided that I could never explain Greeks to management. And I thought if I can't explain Greeks, I can't do this anymore. And so, I went and started consulting and then one day I woke up and I realized that I didn't need to explain Greeks in order to build option portfolios. We needed to know Greeks in order to price them.

You can hedge just like you can construct any portfolio, because at the end of the day, constructing a portfolio, a hedge portfolio, is taking out the correlation of your liability to whatever factors you can identify. That's all it really is. You're taking what is effectively a box of snakes and turning it into cash.

When you have a box of snakes, you can't hold all the snakes still simultaneously. But you can keep them in a box. You can keep them bounded and if one of them gets out, you put that box in a bigger box and then they're still in the box. And it turns out that you can extend that methodology and look at the effective hedge portfolios on real-world financial metrics, whether that's income, claims, capital, whatever. You can generate functions of the market that you know your company cares about and then you can hedge against those and what may or may not be a sufficient way, that is at least statistically defensible. It generates unique portfolios usually. It is not permanent. It can't be because we don't know enough about the future frequently to buy all the instruments that would be required. If we knew all about the future, at least in the statistical sense, we might find that the instruments are not available. So this is how you do it.

You can calculate Greeks, that's one way to do it. Generate correlated realizations of boundary variables. That's a pretty geekish way of saying, model the real world, so you say the real world one month out is going to look like something. It's going to have option prices, interest rates and equity levels, and I think those go together in some sort of statistical way. I can generate simulation of that outcome. And I can value options under that outcome.

All you need to do to construct a hedge is to decide on what you want to avoid by hedging. So if you want to avoid model error tracking that is generated by the change in option valuation, you would use boundaries like the items that go into valuing options. If there are other things that you want to avoid, like GAAP income surprises or statutory surplus impairment, those are the choices.

There's a lot of flexibility in this approach. And more to the point, you don't have to use the words delta, vega or gamma to your CFO, even if he understands them.

You know, the insurance industry really has not made the full transition from the law of large numbers to a financial engineering framework. And insurance will always be about the law of large numbers. The transition will be in what cases that applies and can it be used effectively, and in what cases there are underlying correlations in the cost of the funds that we raise and invest. That requires some sort of action to prevent calamity or to guarantee the ability of a company to carry out its business plan over a period of several years. Those are all legitimate reasons for hedging. It's not about matching the Greeks.

Companies care about a lot of different things. Good hedging programs meet a lot of different needs. It is so hard to get a handle on how you calculate Greeks. How do you value options? How do you run things fast enough? What do we need to know in order to make this a production activity? I think we're still in the early days. What the industry will be using in five to 10 years is, maybe being thought about now, and maybe it'll come about because people like those I see here today ask the questions over and over again. We're thinking about it. I think our clients are thinking about it, and there's a lot of work to do.