RECORD, Volume 28, No. 3*

Boston Annual Meeting
October 27–30, 2002

Session 121PD
Risk Management Issues for Variable and Equity-Indexed Annuities

Track: Investment/Financial Reporting

Moderator: CHARLES L. GILBERT
Panelists: CHARLES L. GILBERT
K. "RAVI" RAVINDRAN†
DARIN G. ZIMMERMAN

Summary: Attendees learn of the risks inherent in common variable and equity indexed annuity designs. The effects of changes in interest rates, market level, implied and actual volatility on economic and accounting results are discussed. Through the use of illustrative examples, this session examines the use of realistic versus risk-neutral scenario sets for valuation, pricing and risk management of these products. Hedging and reinsurance approaches and their impacts are reviewed.

MR. CHARLES L. GILBERT: Welcome to Session 121, Risk Management Issues for Variable and Equity Indexed Annuities.

Equity market experience over the last years has created a lot of interest in this topic. We've all heard of the large write-off that one reinsurer was forced to take recently after choosing not to hedge its equity risk. But there have also been a number of success stories about companies who weathered the storm by at least partially hedging their exposure. At a session earlier this week, I heard someone rhetorically ask where the actuaries were when the equity market bubble peaked in 2000, and comment that insurers who wrote guaranteed minimum death benefits (GMDBs) were, for the most part, running the risk naked. Although there were a number of actuaries out there trying to deliver their message to companies, in many cases there was some unease with respect to using derivatives. There was also a long education process required to sell non-reinsurance risk management solutions internally.

---

†K. "Ravi" Ravindran, Ph.D., not a member of the sponsoring organizations, is president of Annuity Systems Inc. in Toronto, Ontario, Canada.

On Monday, the risk management task force held a well-attended session, "Risk Management Buzz Group, Session 44BG." I sat in on the equity risk topic, where there was some good discussion. One of the issues debated was using models to measure the equity risk and when it is appropriate to use risk-neutral versus risk-real-world scenarios. Another discussion centered on the number of companies still trying to get a good handle on how to properly measure their risk exposure with their variable annuity and equity indexed products.

We have three speakers this morning. My name is Charles Gilbert. I'm an independent consultant based in Toronto specializing in risk management. In 1998, I co-authored a Canadian Institute of Actuaries (CIA)-sponsored research paper on equity risk associated with segregated fund guarantees. I was part of the organizing committee that put together the symposium on stochastic modeling for variable annuity and segregated fund guarantees, held in Toronto in 1999. Over the last four years I've worked with companies in the United States, Europe and Canada on equity risk management issues.

Darin Zimmerman is assistant vice-president, actuarial, at Americo Financial Life and Annuity Insurance Company in Kansas City. His main functional responsibility is modeling, which includes the cash flow testing model, the annual planning model, the liability duration analysis model and financial accounting standard (FAS) 97 and 133 modeling. He has been heavily involved in Americo equity indexed annuity (EIA) management. Americo wrote its first EIA in the fall of 1998 and has grown to become one of the top ten writers of EIA in the United States. Prior to joining Americo, Darin was a consultant with Tillinghast-Towers Perrin in Atlanta for three years. Prior to that, he worked at Minnesota Life, formerly Minnesota Mutual, in financial reporting and variable life product development.

Dr. Ravindran is the founding partner of Annuity Systems Inc. Being the pioneer of applying holistic risk management approach to dynamical risk management exposures relating to variable annuity and equity indexed annuity business, he is one of the few professionals who has the rare blend of exotic derivatives portfolio management experience in nearly every asset class. In terms of variable annuity and equity-indexed, annuity-related experience, he has personally managed risks for portfolios in excess of $80 billion, and consulted for companies whose portfolios totaled in excess of $200 billion. Dr. Ravindran's qualifications include past adjunct professorships in Canadian universities, authoring the well received book "Customized Derivatives," a step-by-step guide to using exotic options, swaps and other customized derivatives, and editing or authoring other books and papers.

We're going to start off by illustrating some of the risk profiles of the different variable annuity and equity-indexed annuity products that are out there. Darin will discuss sensitivity analysis for an EIA case study. I'll talk about some of the real-world versus risk-neutral modeling considerations. Ravi will finish by talking about different risk management approaches.

As most of you already know, the expected cost with variable annuities and GMDB products, in particular, is almost trivial; it's very small. The issue with these products is that the risk tends to be fat tailed. We're going to look at the risk profile of three variable annuity products. We'll see that the risk exposure is a function of the product features. In the risk profiles, what we're looking at is not just the risk associated with the payoffs under guarantees; we're also interested in the revenue loss associated with the lower fund balances. Often times on GMDB products, this is the greater risk exposure for companies.

We'll look at three GMDB products. We'll look at a flat return of premium guarantee. Then we'll look at a GMDB that has an annual ratchet, and then we'll add a 5 percent rollup and take a look at what the cost looks like on those products. Just to give you an idea of the fair risk charges, you can see that for the flat guarantee it's not very big, only eight basis points. But once you add the annual ratchet, it increases significantly to 46 basis points, and then adding the 5 percent rollup on top of that is 70 basis points.

Figure 1 shows the risk profile for a flat GMDB product. Take note of the scale that we're using here. If you look at the curve measuring the profit, at the extreme left tail it dips to 60,000, and at the extreme right it peaks at 20,000. Because the tail exposure is very sensitive to the number of scenarios that you choose in your modeling, we're more interested in looking at something like a conditional tail expectation (CTE), which companies in Canada tend to look generally at a CTE (95 percent). So for this product, the CTE (95 percent) for the flat GMDB is approximately negative 33,000 or 34,000.

Figure 1

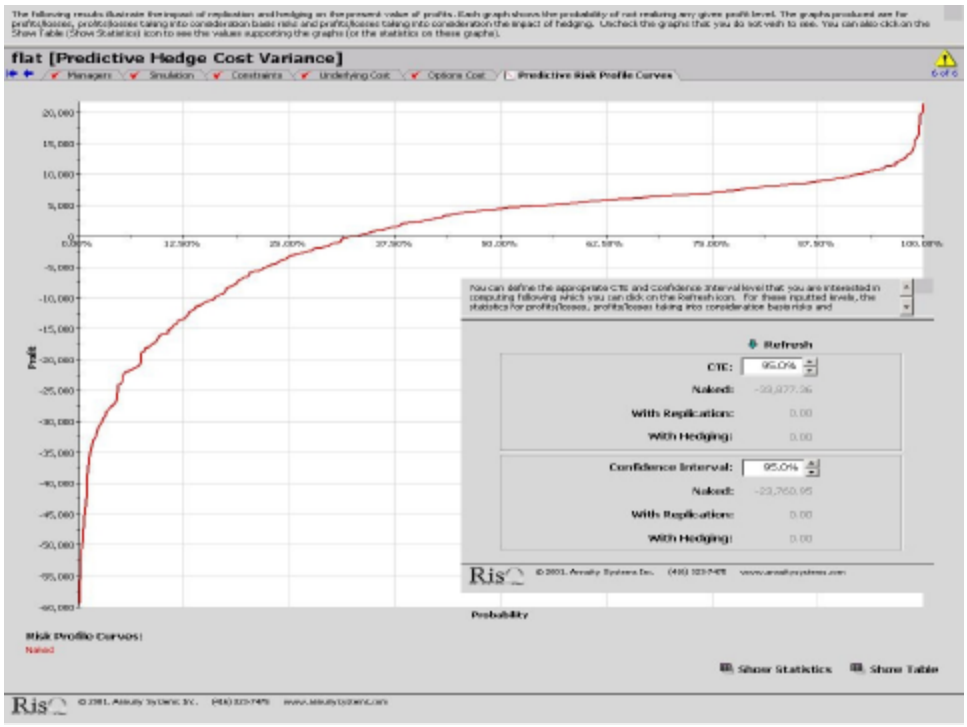

In contrast, if you add an annual ratchet feature to this product, take note of the scale on the extreme left in Figure 2. The profit drops to negative 275,000. The more significant or meaningful measure to look at is the CTE (95 percent); it's gone from negative 34 to negative 61. If we add a 5 percent rollup, we see as shown in Figure 3 that the CTE (95 percent) goes from negative 61 to negative 78—a pretty big increase. I also want to point out that the extreme left tail exposure on the ratchet and rollup product drops to negative 175,000, or almost 200,000, whereas on Figure 2, the extreme exposure dropped to negative 275,000. That's purely a function of the number of scenarios that were run. For this example, we ran 5,000 scenarios. We could have run significantly more, but wanted to illustrate the sensitivity of using a lower number of scenarios, and how that can sometimes give you a counterintuitive result if you're focusing only on the left tail exposure.

Figure 2

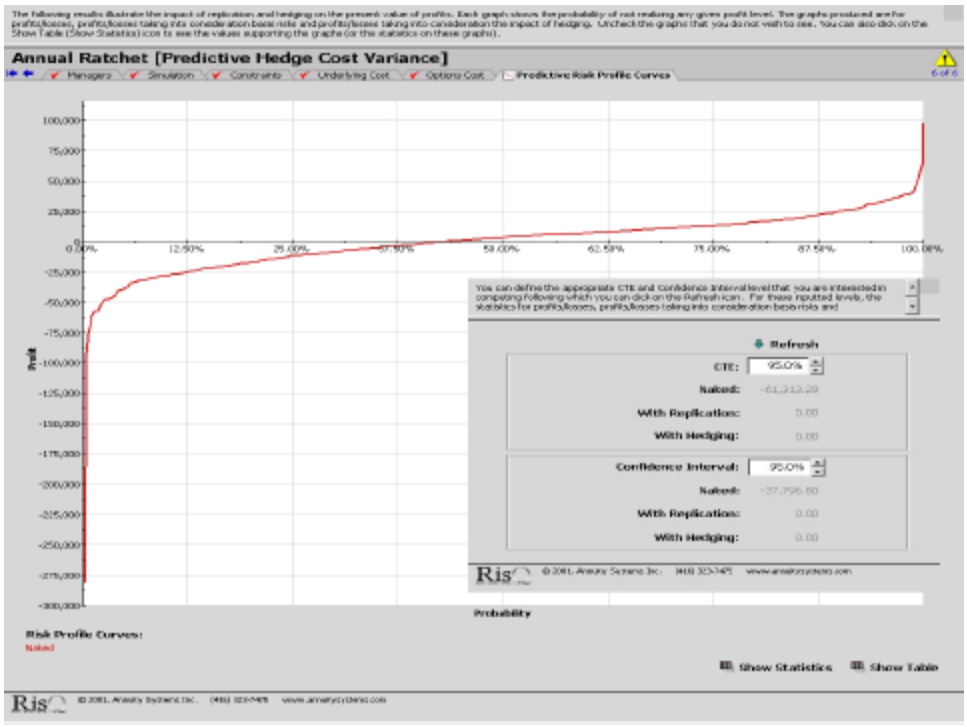
Annual Ratchet [Predictive Hedge Cost Variance]
Managers
Simulation
Constraints
Underlying Cost
Options Cost
Predictive Risk Profile Curves
6 of 6
100,000
75,000
50,000
25,000
0
-25,000
-50,000
-75,000
-100,000
-125,000
-150,000
-175,000
-200,000
-225,000
-250,000
-275,000
-300,000
0.00%
12.50%
25.00%
37.50%
50.00%
62.58%
75.00%
87.58%
100.08%
Profit
Probability
Refresh
CTE:
95.0%
Naked:
-61,212.29
With Replication:
0.00
With Hedging:
0.00
Confidence Interval:
95.0%
Naked:
-37,796.60
With Replication:
0.00
With Hedging:
0.00
Risk Profile Curves:
Naked
Show Statistics
Show Table

Figure 3

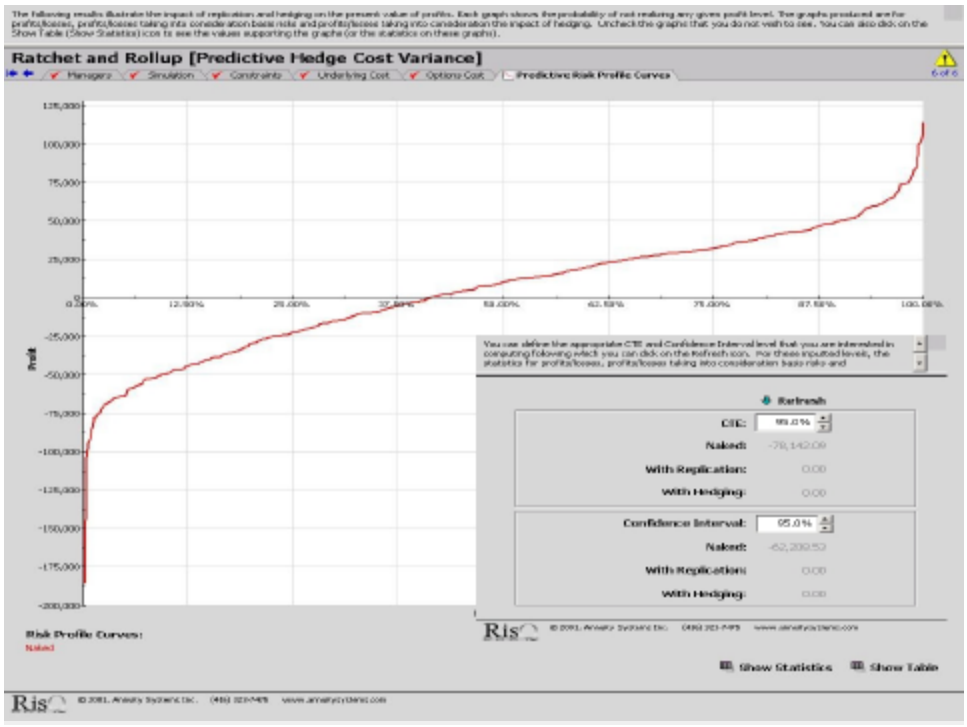

I'll let Darin speak about the equity-indexed annuity products.

MR. DARIN G. ZIMMERMAN: Thank you, Charles. The risk inherent in the equity-indexed annuity depends on the underlying assets. The analysis on the equity index slides is a little bit different from the analysis for the variable annuity features, because we're talking about a product versus features. The product has to be hedged and the premium hedged, whereas with the variable annuity it's just the premium invested or spent for the benefit, and then what is the ultimate cost of the benefit? Figure 4 shows my assumptions, slightly more than there are with the variable annuities. I'll let you in a little secret: by far the worst assumption on earth, the fifth from the bottom, is the participation rate fixed at 50 percent, annual ratchet.

Figure 4

Assumptions for sample EIA

Age	35
Yield Curve	9/30/01
Implied Equity Growth	9.8%
Volatility	19.2%
Lapse	none
Mortality	none
Investment spreads	NAIC1 bonds
RBC	200% of formula
GMV	3% of 75% of FYP
Surrender charges	15%, 14%, … 1%, 0%
Participation rate	Fixed at 50%, annual ratchet
Commission	15%
Acquisition expenses	$50/ policy + 1% of Premium
Maintenance expenses	$30/pol @3% infl + 10 bps
Average size	$100,000

That's a bad assumption because these products generally won't guarantee the participation rate. If you manage your equity index block like you manage a deferred annuity block, you are going to have an option budget. You're going to solve for the participation rate with risk-free rate, volatility at the time, the design of the option, point-to-point averaging, binary duration to maturity, and you're going to change that number. Being able to do that in the model takes a little bit of programming savvy, so for this exercise we left it fixed at 50 percent, which is probably a bad assumption. But it does give us some insight into the risk profile of the product.

In Figure 5, we see that the risk profile is fairly fat tailed. This is a discounting of the distributable earnings including taxes and risk-based capital (RBC) at a 10 percent discount rate. If you're pricing with a 10 percent hurdle rate, you'll be interested to see where it crosses zero. This scenario assumes the investments strategy is 100 percent bonds, which is also know as running the risk naked because you are completely exposed to movements in the underlying index. This is what the risk looks like for writing this particular annual ratchet equity-indexed annuity. Again, you see it's fat tailed, and you have to remember that things are kind of opposite. You lose money when the Standard and Poor's (S&P) or your index goes up, because you're running the risk naked. When the S&P doesn't go anywhere, that's when you make your money.

Figure 5

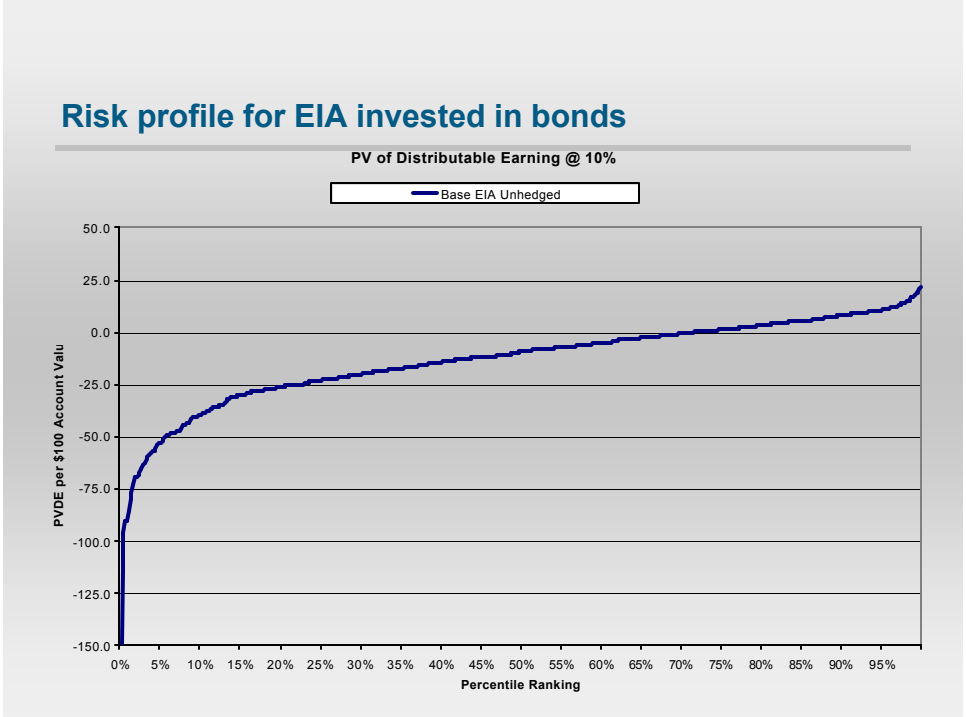

Figure 6 shows a couple of different investment strategies. The blue line is the base run from the previous figure; it's all bonds with cash flow matching. The red and brown lines are invest in a combination of bonds and the underlying index. If you are trying to match the 50 percent participation rates and you actually buy the index, you'll need 70 percent bonds, because at that level you have roughly 20 percent invested to cover your RBC. So a $1,000 fund is going to need $1,200 of assets, given the current RBC guidelines. With the brown line you're pretty much over-hedged, investing 80 percent in the index with 50 percent in bonds. Because you have more equities, you have more RBC and need roughly 130 percent of the fund amount in assets. The pink dots and the green dots show two particular scenarios as examples. Scenario 129 is the pink dot, which had very high equity returns. Scenario 382 had low equity returns; it's the green dot. Here's how the dots work: If you look on the blue line you will see the pink dot close to the left-hand side. Since the index rises substantially and the risk is not hedged, the company loses a lot of money. It looks like roughly 4 percent of the scenarios caused the company to lose more money than scenario 129.

Over on the right, the green dot touches the blue line almost at the right edge. Since the index didn't go up, the company's liability didn't go up, but they got to keep all of the bond investment income. So the company made money. It looks like less than 1 percent of the scenarios caused the company to make more money than scenario 328. On the brown line, where the company is over-hedged, the position

of the dots is reversed. The company makes money when the index goes up and loses money when the index goes down. It is interesting to note that when the company is closely hedged, one of the extreme scenarios falls right in the middle.

Figure 6

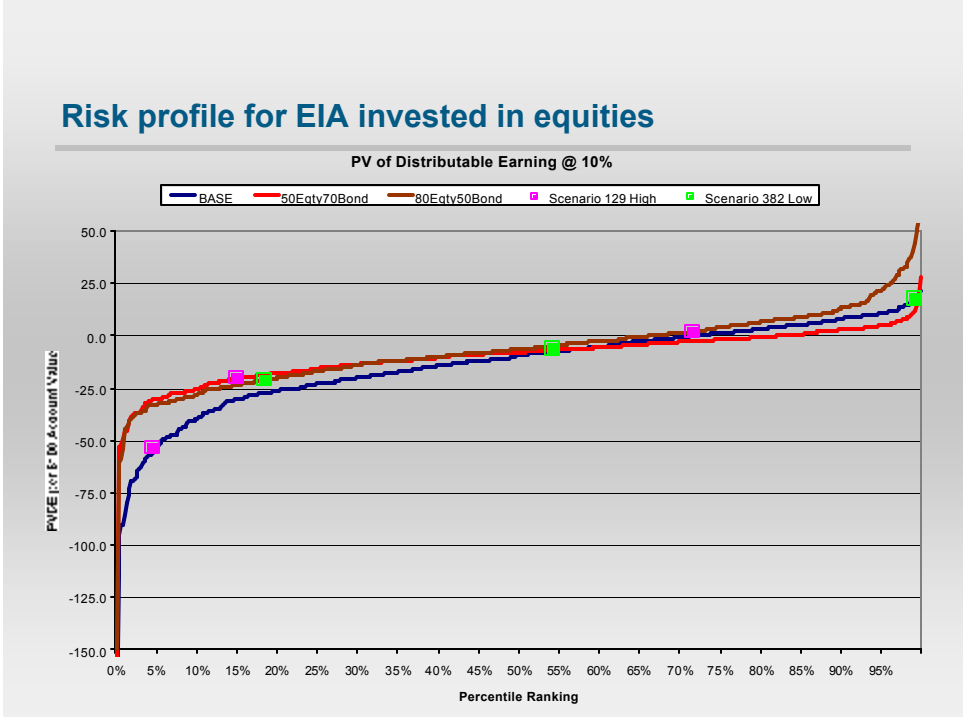

The sensitivity analysis on Figure 7 shows you how the options change as some of our assumptions change. We're going to look at changing the level of interest rates, we're going to look at changing the implied growth rate, so our scenario generator is your typical lognormal scenario generator. It has a drift in the neighborhood of 8 percent, so we made that 8.8 percent, increased it 10 percent, and with volatility we raised our volatility assumption by 500 basis points. So it went from 19.2 to 24.2.

Figure 7

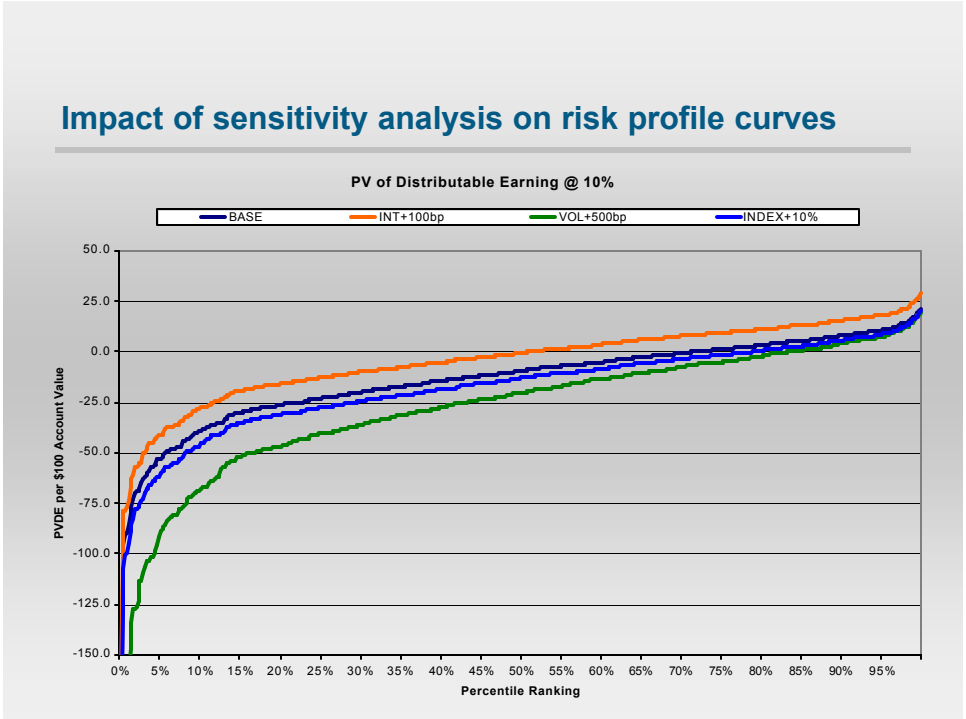

If the orange line is higher interest rates, if you're writing this product in a higher interest rate environment and you can get away with still offering a 50 percent participation rate, you're going to make more money whether you hedge it or not. Likewise, if there's more volatility or if the implied growth rate or if the S&P grows more, then under the base model you're going to lose more money, because again, this is all running the risk naked. It's not hedged.

Figure 8 is the picture in my head of what a good dynamic lapse rate formula would look like for equity-indexed annuities. This isn't necessarily for the annual ratchet. There are some annuity designs out there that are 10-year point-to-points. If you surrender, you get the guaranteed asset and there are some vesting schedules, but your ordinary dynamic lapse formulas aren't going to work. They look at a credited rate, and you haven't credited anything. Until the tenth year, they are only getting their guaranteed return. So for a 10-year point-to-point, the rational behavior isn't going to be a function of credited rate, it's going to be a function of the current and S&P rate and distance to maturity. What I've tried to capture is that, on the far part of the scale where it looks like a cliff, policyholder behavior is going to get very rational. If you're six months away from cashing in and the S&P needs to go up 50 percent in order for you to get your $1,000, you can get your $1,000 today. So a lot of people are going to do that. Likewise, if you're six months out and you're 50 percent in the money, nobody is going to take their money out, so you're going to be down at the zero.

Figure 8

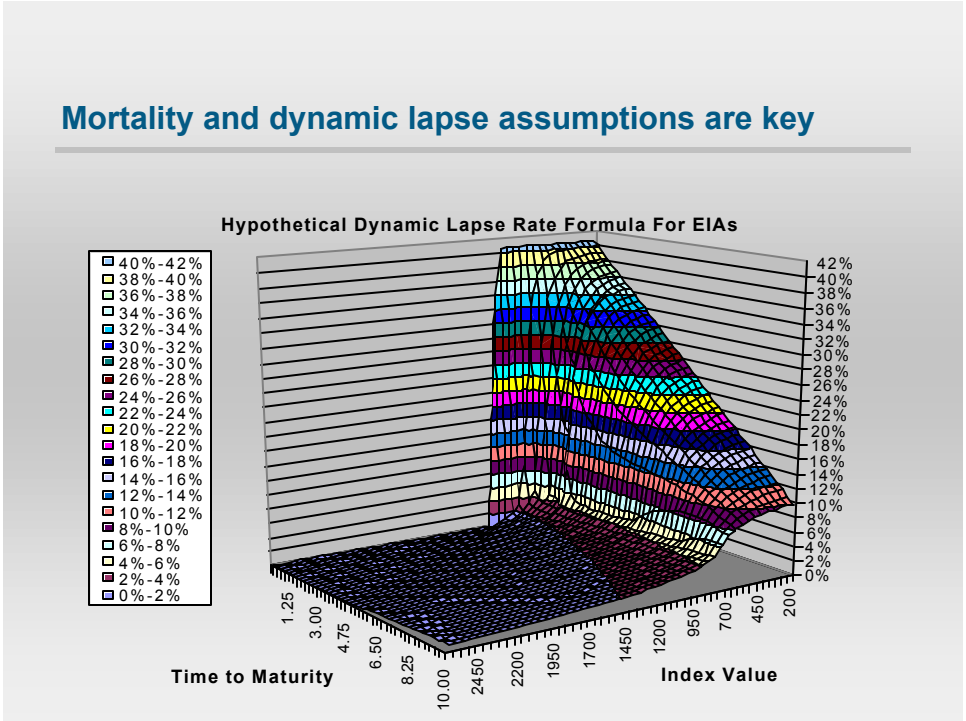

Before we get into what the accounting results look like when they are hedged, the question is what are you going to hedge? Are you going to hedge your statutory accounting? Your GAAP accounting? Your economic value added? The problem here is you've got to pick one. You can't hedge all three. All of these accounting methods have different assumptions, have different methods and are going to give you different results. You're going to pick one and have volatility in the other two. So whatever you're trying to manage to is probably what you're going to pick to hedge. But be aware you're going to have volatility in the others.

Figure 9 shows the hedge results. The three lines at the bottom are the three graphs we saw before. The blue is the base, the brown is the 80 percent equities, and the red is the 70 percent bonds and 50 percent equities. The light blue, orange and purple are the hedge results, showing how you can immunize your profits from volatility in earnings. The purple hedge is the stat reserve, modeled here as the fund less the surrender charge, just for simplicity. The blue line hedges the entire fund, which would be akin to GAAP accounting before FAS 133. Then with FAS133, it's a total crapshoot. We have no idea what it is, so I randomly picked hedging 150 percent of the result, just to illustrate the base assumption in the model, which is that equities are generally a better bet than bonds. So you look at this and think that if you hedge 150 percent, you're going to make tons of money, so we should be long the index. But again, there's the underlying assumption that's just a bias, that equities are going to do better than bonds. That is a good assumption, but are you willing to bet your company on it? Most people are going to pick the blue or the

purple line, because they don't want to take an equity position only based on the assumption that since equities were better in the past, they're going to be better in the future. Pity the fool who made this recommendation two years ago.

Figure 9

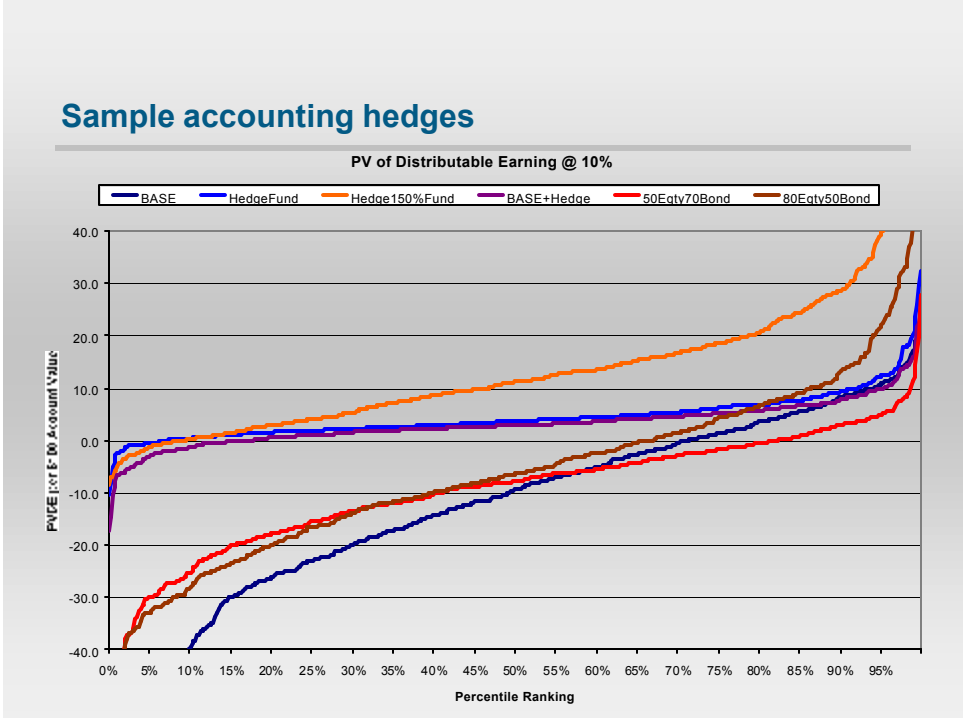

Figure 10 is a picture of what you're actually trying to hedge. If you're trying to hedge the economic value added, this is the shape of that option. This is back to the 10-year point-to-point option that I picked here just because it's a better illustration. You could just as easily shrink this down to one year, but it would be a lot thinner; this gives a better picture of it.

Figure 10

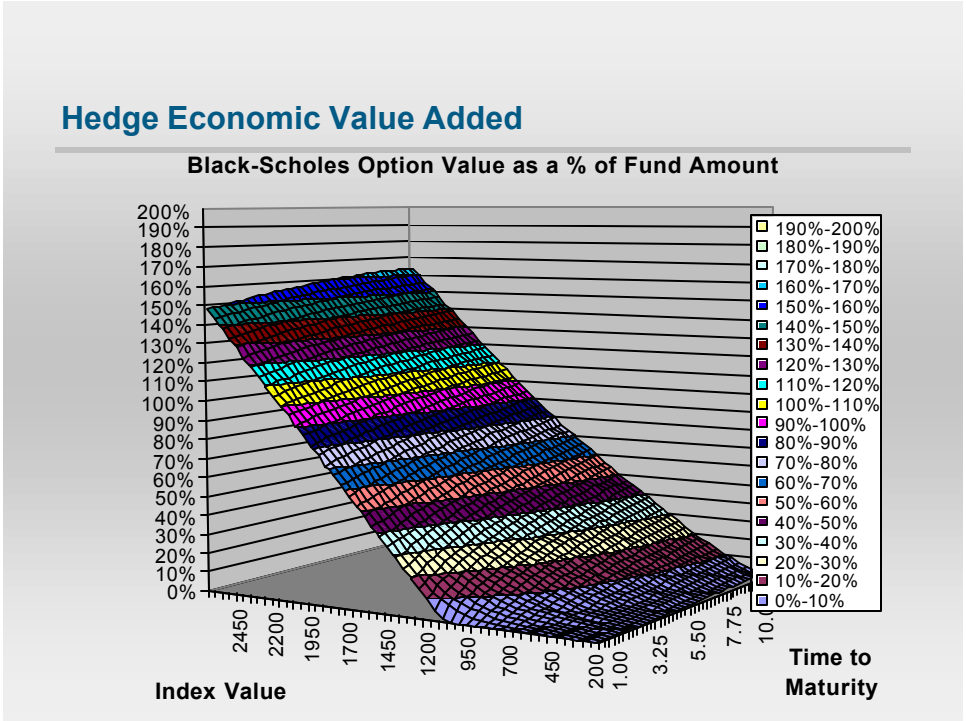

If you were hedging stat income, the ramp that comes down on the right side of Figure 11 is basically your cash value floor. Underlying this assumption, you have some bonds and you've purchased an option. The sum of those bonds in that option value always has to equal the cash value. Those bonds grow with interest in time. So as you come forward on the left axis, as time gets shorter your bonds grow, so you can allow your option value to go lower and still be able to pay the cash value without taking the loss.

Figure 11

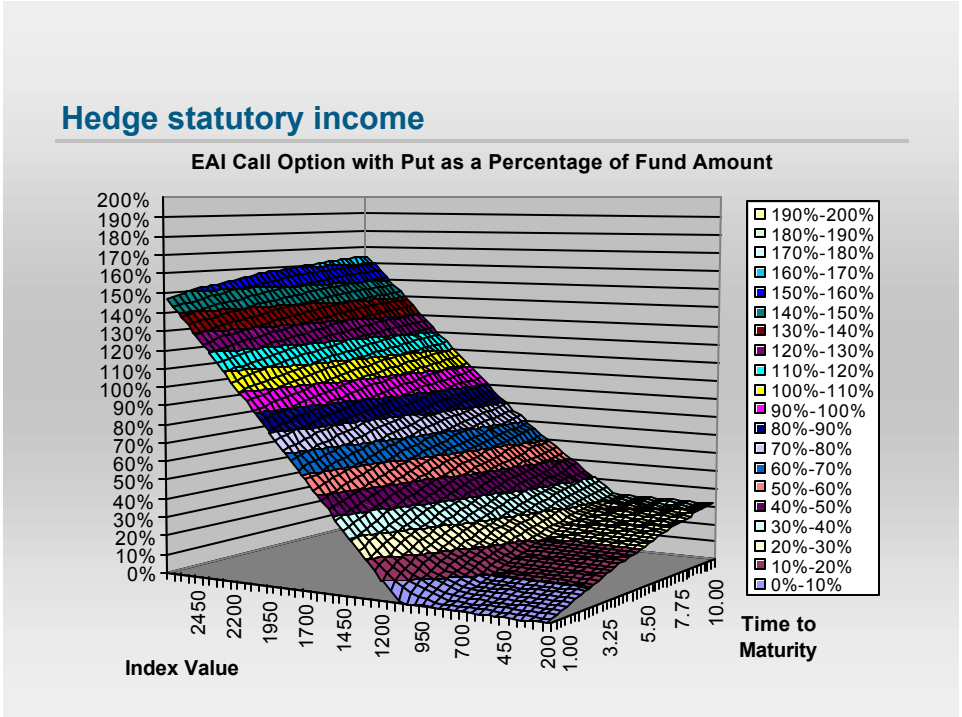

At Americo, we got away from a reinsurance approach. Now we're managing the risk with a dynamic hedge portfolio. Figure 12 shows the "Greeks." For anybody who is considering bringing their management of equity index in-house, these are what you're going to live and die by. Just in case anybody doesn't know, the Greeks are the price sensitivities of your portfolio with respect to either your liabilities or your assets. Delta is the change in price with respect to the index. So as the S&P or the Dow goes up or down, each point will change the value of our liabilities by $39,000; each point will change the value of our assets by $36,000. So our Delta is approximately $3,000. So as the S& P moves one point, our value changes by $3,000. Gamma is the second derivative of price with respect to index. As it moves one point it's 39,375; as it moves two points, that second point is 39,375 plus 204. It's the convexity of it. Then the other ones are a little bit simpler than that. Vega is the change with volatility. If volatility moved up by 100 basis points, our liabilities would go up by 500,000. Rho is interest rate, and Theta is the time decay, but we don't usually manage to that.

Figure 12

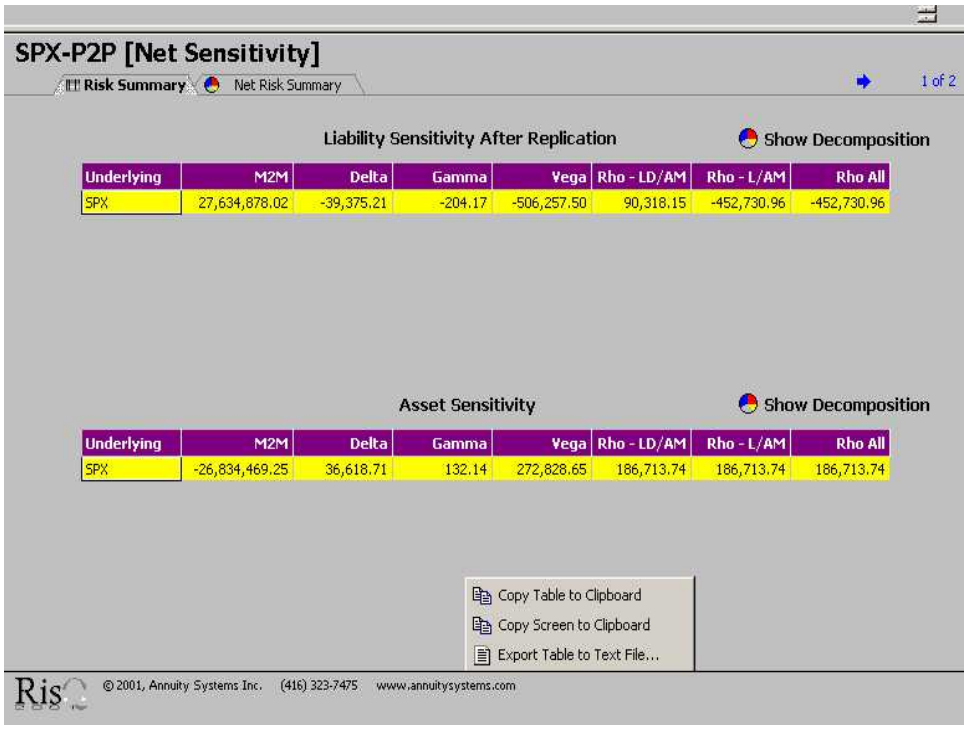

Another thing to be aware of when you're managing is that there's a skew to the volatility curve in the marketplace. If you price marketable options, the implied volatility will look something like Figure 13. The reason to be aware of this is that this can hurt you. If you've got a portfolio of different products—point-to-point options, binary options, one-year maturities, two-year maturities, 10-year maturities, different participation rates, different caps, different lift-offs—you could have liabilities that are spread out all over this surface. If you're trying to manage this with futures and several big over-the-counter options, if you only go to the market once a month and only buy great big options, as this surface flexes up and down you're going to have a mismatch between your assets and your liabilities. That can create earnings volatility, and management never likes earnings volatility. So this is another thing that people need to watch out for. You set it up carefully with futures and options, and then there's a big flex, all of a sudden you've lost $10 million dollars. What happened? Well, you forgot there was skewed volatility.

Figure 13

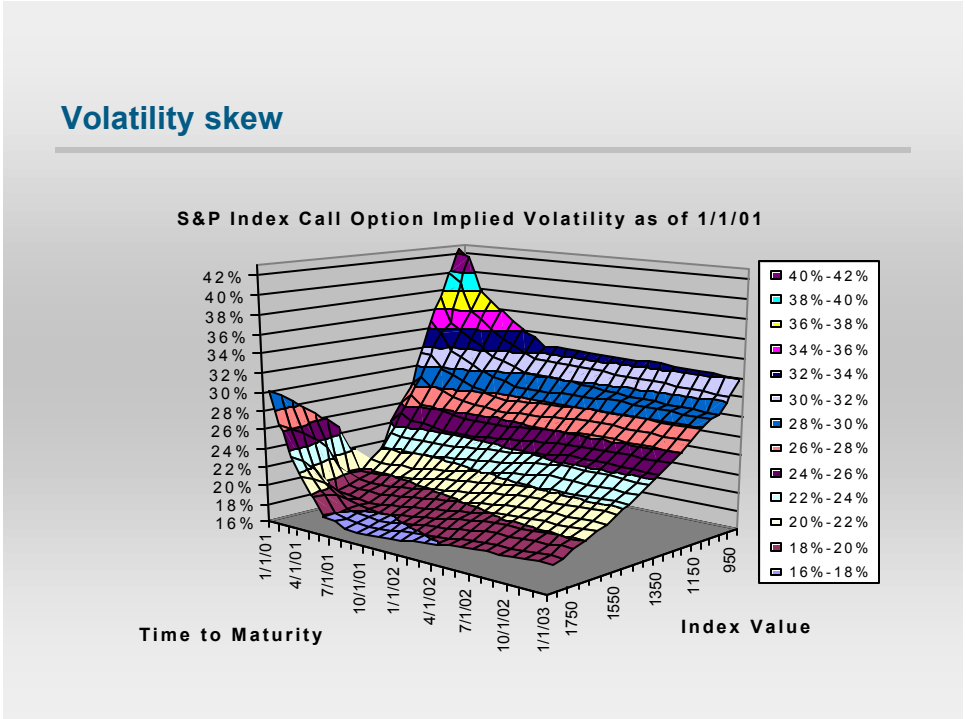

Figures 14, 15 and 16 illustrate that constant growth projections for ratchet equity-indexed annuities are really bad news. When we do the cash flow testing at Americo, we run stochastic scenarios. For our plan, we run stochastic scenarios and mash them all together. The reason is you can't pick one scenario and have it adequately represent an equity-indexed annuity, because most people are going to pick a constant growth scenario. But what people forget is any number you pick in a constant growth scenario is mathematically proven guaranteed to give you the smallest account value. That's what Figure 14 shows. This shows a two-year ratchet, instead of an annual ratchet, so there are only two maturity periods for a four-year period. If we have constant 10 percent growth in the index, the black line in the middle gives you a fund value of $1,215.51. If you change that just a little bit, either you grow a little faster in the first two years, or grow a little slower in the first two years, then you end up at the exact same point with $1,215. 53. With the orange and the yellow lines, if you grow significantly faster and then significantly slower and end up in the same place, you get $1,216.88, and if you go up and then down, or down and then up, $1,237.00.

Figure 14

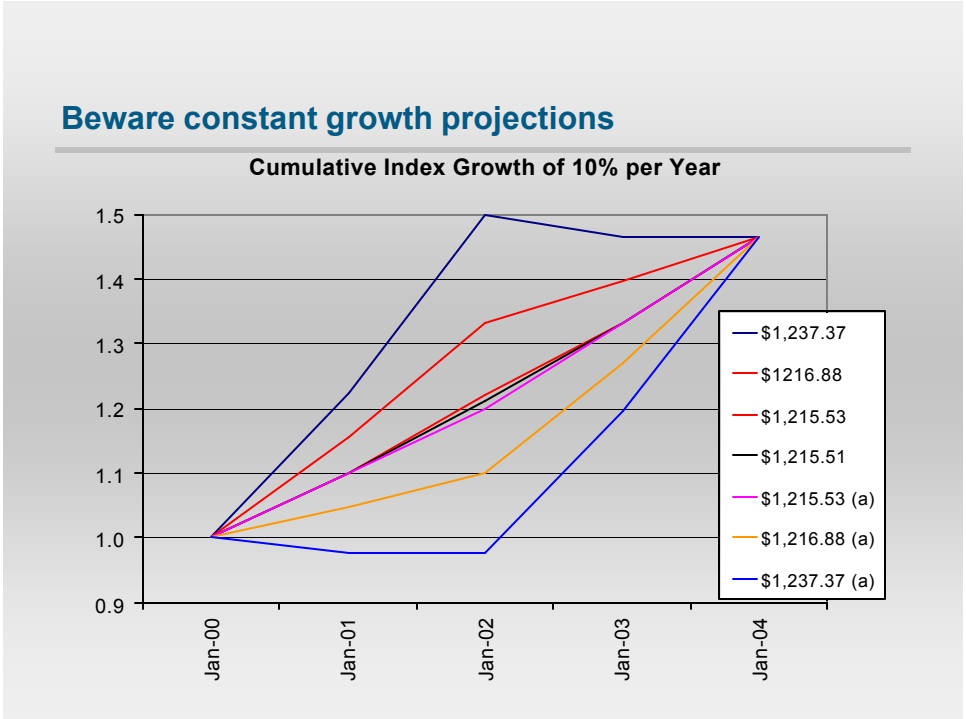

Figure 15

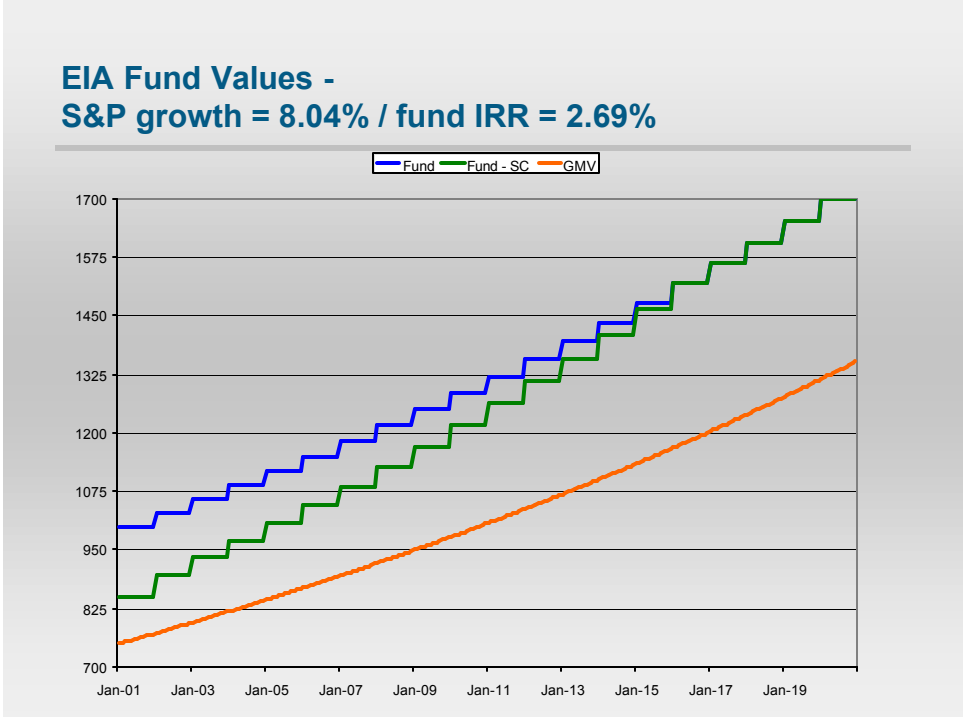

Figure 16

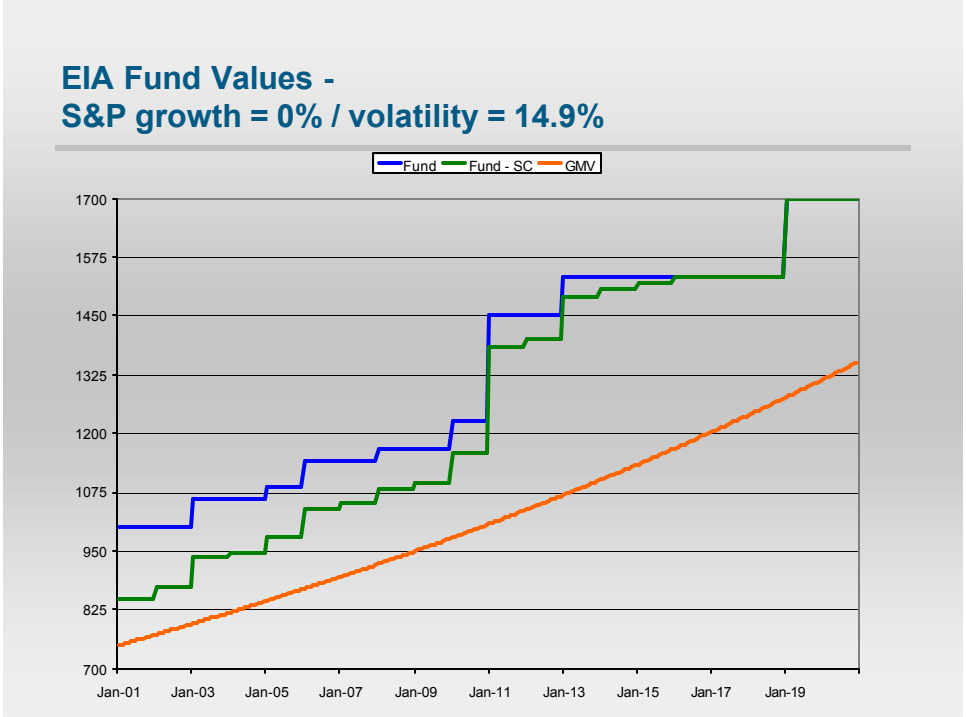

The implications of this: Management asks you, "I want to assume an 8 percent growth rate. What's this product going to do?" As Figure 15 shows, you plug it in, it grows at 8 percent. The ending fund on a $1,000 deposit after 20 years is $1,700. Your president will say, "Why are they only getting a 2.69 percent credited rate on 8 percent S&P growth? Nobody is going to buy this; this is a terrible product." Well, I started playing with some randomly generated lognormal equity S&P scenarios and I got what you see in Figure 16. I got a scenario that started at $1,000 and ended at $1,012. There was 0 percent index growth over the period, but my fund still grew to $17,000, because of the interaction of the ratchets and the resets. So when your president says, "You mean if the S&P grows 8 percent, the client is only going to get 2.69 percent fund growth?" you can tell him or her, "No, if there's 8 percent growth, that's the guaranteed minimum." In the highly unlikely scenario that it's 8 percent growth with zero volatility, then that's what they'll get. The actual expected number is going to be some function of the implied growth and the volatility. This is just the smallest.

MR. GILBERT: I'm going to talk about real-world versus risk-neutral modeling issues. These two approaches really provide two different perspectives. When you're looking at the guarantees in variable annuity and equity-indexed annuities, these represent embedded options whose paths will be a function of the underlying equity prices. One of the challenges for the actuary is choosing which approach to

use for which given circumstance, whether it's valuation, pricing or risk management of the guarantees.

Stochastic simulation of future equity returns is generally used with either real-world or risk-neutral probabilities. Stochastic simulation is not always necessary under the risk-neutral approach. There will be times where a closed form solution to the problem exists, in which case you can use a formulaic approach such as Black-Scholes. Or there will be times when you can use the manufacturing cost or derive the cost of the option by using available tradable securities that are priced in the capital markets and replicating the payoffs of the option. Under those two circumstances, you don't have to use stochastic simulation to derive the price.

The risk-neutral approach provides a fair market cost, whereas the real-world approach provides an expected or best estimate cost. In terms of the terminology, the real-world approach is referred to as the traditional actuarial approach, the P-measure approach or equilibrium pricing. Real-world simulation incorporates a real world view. For example, many companies will look at historical equity returns and say, "Well, we believe that they will continue to earn an expected rate of somewhere around 10 percent." Best estimate assumptions are based on historical returns. Under the Geometric Brownian Motion model, a real drift is used. Real probabilities are used. When you are doing your stochastic model, if you're running 10,000 scenarios, the sample paths are equally weighted. The one big advantage of using the real-world approach is that you can reflect your actual assets and your reinvestment strategy. It's recognized by a number of academics and the Canadian Institute of Actuaries as providing the best estimate for the liabilities. In terms of trying to measure the risk, it provides the risk profile curve, confidence levels and conditional tail expectation (CTE) measures.

The risk-neutral approach is used to determine the fair market price of the option. It's also known as the capital markets approach, the Q measure approach or risk-neutral valuation. Risk-neutral simulation incorporates the capital markets view. That means that the only rate of return you can earn with certainty is the risk-free rate. It assumes that investors are indifferent toward risk; the risk-neutral drift is adjusted. One of the criticisms of the risk-neutral approach is that it uses unrealistic probabilities. When you do a risk-neutral valuation, the sample paths are risk adjusted using risk-neutral probabilities. Risk neutral approach is arbitrage free and produces the fair market price of the option, i.e., the manufacturing cost.

I want to highlight three steps to risk-neutral valuation, specifically using a binomial model. The first step is to solve for the risk-neutral probabilities. The second step is to calculate the expected payoffs for the option along your binomial tree. Then, step three, discount those back using your risk-free rate. Solving for risk-neutral probabilities in step one is mathematically equivalent to assuming that the expected return is equal to the risk-free rate. That gives you a bit of a bridge between the risk-neutral world and the real world.

So when is it appropriate to use each approach? The real-world approach is generally used for valuation and provides the best estimate of the policy liabilities. It's the recommended approach in Canada for valuation of segregated fund guarantees. In Canada, the reserve is based on the CTE. The risk-neutral approach is generally used for risk management. There is some debate about whether or not the real-world approach should be used for risk management. I think that when you want to look at your risk exposure, you want to look at your risk profile. You want a best-estimate view of that. From a risk management perspective, what would it actually cost if you were to hedge your exposure? So it's the actual cost of the risk management strategy that would be available in the capital markets. There is an argument that a real-world approach could be used if you're using tactical hedging strategies that incorporate a view on the markets.

It's not totally clear what should be used for pricing. I believe that it depends on the risk management strategy. If your intention is to hedge the guarantee, then I can see an argument for using a risk-neutral approach. If, on the other hand, you're planning to run the risk naked, then you may want to use the real-world approach to get the expected cost of that. I know of at least one large multinational insurer that developed pricing standards that base the price on the real-world approach, but only after doing the analysis and looking at what the cost would be on a risk-neutral basis, and taking that into consideration.

Another approach is to explicitly reflect the risk management strategy in a stochastic model. Basically what we're talking about here is stochastic on stochastic simulation. If you're going to project 1,000 or 10,000 paths, at each model point, at each node for each of those paths, you would calculate the cost of hedging. You can quickly get into literally billions of calculations for one product, for one policy. It's computationally very time intense.

People in academia are quite adamant that the capital market approach should be used for pricing equity-indexed annuities or variable annuities, or anything that has an embedded option in it, because this is how the capital markets price options. Since the guarantees are embedded options, this suggests that they should be priced based on the manufacturing cost, which equals the price of the tradable securities that replicate the option payoffs, i.e., no-arbitrage. People who subscribe to this view suggest that any other price is the wrong price for two reasons. One is that arbitrage opportunities will exist if your price is lower than the manufacturing cost; people could go out and replicate the paths and arbitrage your company. The other is that if your price is higher than the actual capital market price, then no one is going to purchase your product or your option.

In reality, most insurers do not price using the capital market approach; the price is generally lower than in the capital markets. We don't see any arbitrage occurring. At least not yet, and one of the conditions for arbitrage to happen is that you have to have complete markets. Complete markets don't exist for mortality and lapses. Equilibrium pricing is used by many companies. Again, that's consistent with the

risk management strategy of running the risk naked, which reflects what most companies are doing currently. Equilibrium pricing can also be used to reflect the hedging strategies that incorporate a view on the market. When you're doing stochastic on stochastic simulation, it may be appropriate to use a real-world approach to calculate the cost of the derivatives in your hedge at each of those model points. I think that would be the case if you were taking a tactical hedging strategy and trying to incorporate a view. The trade-off there is you're not getting a perfect hedge; you're not hedging the entire risk. By accepting some deviation from that, you don't have a certain return.

Figure 17 illustrates an example for an equity-indexed annuity. This is a 10-year EIA with a 75 percent participation rate. In this case, the insurer has invested in 10-year bonds. The example assumes a 5 percent return. What we've labeled here as the net cost or cost is actually the difference between the return that would be earned on those 5 percent bonds over 10 years, and the expected payoff of the call option using both a real-world approach and a risk-neutral approach. When you're dealing with call options, the relationship is reversed: the cost of the option is more expensive using a real world approach. In contrast, if you're valuing a put option, the cost is less expensive using a real-world approach, because in a risk-neutral world, the expected return is effectively equal to the risk-free rate. So on a call option, if you're assuming the expected return is the risk-free rate, your payoffs under the call option are going to be lower. So for this product, the solid red line shows that the difference between what you're going to earn on your bonds and the cost of the option is going to be 40 basis points. In this picture, it's a negative cost or a profit of 40 basis points, but under a real-world approach the cost of the call option is greater. You actually have a loss of 11 basis points. You can see the risk profile curves and see that the risk-neutral curve is higher at all points, as you would expect, because you're using an expected return assumption of the risk-free rate as opposed to the 9.5 percent we used in this example.

Figure 17

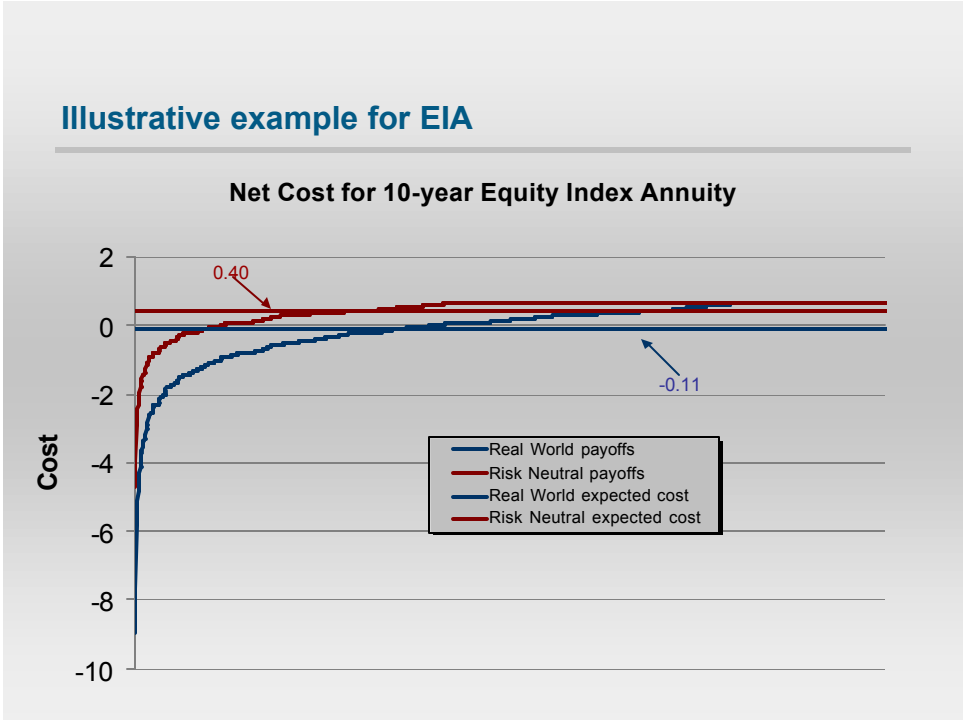

I'll turn it over to Ravi, who will talk about hedging and reinsurance approaches.

MR. K. "RAVI" RAVINDRAN: Thank you, Charles. I want to focus primarily on the different types of hedging strategies that exist in the marketplace today. It doesn't matter whether you're an EIA writer or a variable annuity writer, the following are the strategies available to you.

- Take on your risks naked without any form of hedging. This is consistent with some of the comments that were made earlier about how and when you apply risk-neutral versus real-world type valuation.
- Statically hedge your risks. This approach can only be implemented for certain types of products, but not the whole gamut of products. Furthermore, the definition of static hedging is a nebulous one, because people have used this term to mean different things. For the purposes of our discussion, by static hedging I mean going to the over-the-counter market and then executing trades like the 10-year/5-year/20-year options with the counter parties being investment banks. In doing trades like this, one has to keep in mind:
    - the minimum size of the transaction that the investment bank would require

  - that investment banks typically will not and do not take on basis risk. More precisely, while they would be comfortable quoting on derivatives on well-established or liquid indices (e.g. S&P 500, NASDAQ 100), they would typically be unwilling to quote on an actively managed equity fund (e.g. Templeton Growth Fund) or even an illiquid index
  - even though you may be able to convince an investment bank to quote on an illiquid index that may be highly correlated to an actively managed fund, there's going to be a lot of illiquidity premium attached to that quote.
  - accounting and regulatory relief many times have been the driver that has pushed a lot of the variable annuity writers into reinsurance. But, of course, things are slowly changing, and over the next couple of years you'll probably see a lot more different initiatives to address those issues as it relates to the use of OTC derivatives to manage such risks.

- Dynamically hedge your risks. Instead of buying something and keeping it aside, you get into more short dated instruments, which could be a combination of futures contract, or short-dated options (e.g. one-month, one-year, two-year options). The modus operanda here is to work your program in a dynamic fashion so that you're constantly looking in the marketplace and rebalancing your risks (delta, gamma, vega, rho) to bring them back on track.

  Dynamic hedging is mostly driven by the liquidity of available hedge instruments. For example, if a particular index or fund is used to support a product, then in implementing a dynamic hedging program, one would look for other highly liquid indices with a large correlation to the funds/indices being replicated. Most market makers that sell these OTC derivatives actually deploy such dynamic hedging programs themselves.

  The big advantage that an investment bank has over any hedger of such risks is that they see flows on all different sides of the market. From the insurer's standpoint, the insurer has only its equity-based liability to hedge, whereas from an investment banking standpoint, there are Companies A, B and C with opposing flows, which kind of helps with the whole execution of the dynamic hedging. In dynamically hedging these risks, although one is hedging against market movements, the basis/lapse/mortality risks are still present. The reason for this arises from the fact that there are no financial instruments in the marketplace that contain basis/lapse/mortality associated characteristics. Because the dynamic hedging program allows one to immediately react to sudden realization of such actuarial factors without much trade unwinding, as a consequence, the amount of savings can be quite large relative to embarking on a static hedging program. Like the static hedging program, the dynamic hedging program is also exposed to the same sort of problems relating to accounting and regulatory relief.

- Reinsurance, if it does exist at a price you want, is probably the best solution. But the problem is that it's not generally available. For example, if you look at the variable annuity markets right now, there's probably only one reinsurer in the marketplace. Since many reinsurers enter and exit the market, you just do not know when the next wave of entrants would be coming in. There are a few reasons why reinsurance is very attractive. Some of these are:
    - Reinsurance typically does not introduce basis risks.
    - Reinsurance is very accounting and regulatory relief friendly.
    - Reinsurance is a concept that can be easily explained to the board and senior management, relative to the dynamic hedging concept. Explaining dynamic hedging, transacting into the futures/options contracts, hedge tolerances, etc., can overwhelm people resulting in them scratching their heads and running out of the room. To compound to this problem, boards many times do claim that they want to know how dynamic hedging works but IN 15 MINUTES, which goes to show a board's desired level of understanding.

- Developing new products is a sexy way of addressing risk management issues. The motivation here is to create new products that have offsetting risk characteristics. For example, if you're writing a variable annuity (e.g. GMDBs with ratchets) then products like EEBs or EIAs would need to be developed to offset some parts of the risks. Before doing this, you need to first thoroughly understand what kind of risks you want to have on your books before starting to think in terms of what kinds of products can be created to offset risks that you no longer want. As properly done, this methodology can also be used to offset some mortality/lapse/basis risks.

The following summarizes all the above-mentioned strategies:

- Going Naked: The advantage to this strategy is that there is no cost to it, because there's nothing to pay for. The disadvantage is that you're totally exposed to risks, which can potentially be a very expensive solution if all hell breaks loose—-like what we are continuing to see in the marketplace.
- Reinsurance: It's a perfect strategy if you can get it at a price that you want. In addition to being potentially expensive, you are also exposed to the counter-party risks of the reinsurers. Having said that, the bigger problem with this strategy is that it is generally not available.
- Static Hedging: You've got some element of risk mitigation. It can be expensive and there's the whole idea about over-hedging or under-hedging given the lapse/mortality typically stays with the direct writer. In using this strategy, you would have to adjust the notional size of trades that you require. But the risk with this is that sometimes your estimations are going to be off given the realizations of lapses and mortalities. When you are under-hedged, you need to go back and buy more options, but face the risk of volatilities spiking up resulting in your paying an extra premium. On the other hand, when you are over-hedged, you can unwind your illiquid

positions incurring a huge bid-offer spread. In addition to this, if at the time of your unwinding, your volatilities drop to a level lower than where they were at the time of purchase, the premiums obtained would be a lot lesser than that spent for the purchase. The static hedging strategy pales in comparison against the reinsurance in which you typically mitigate your lapse/mortality/market risks.

- Dynamic Hedging: With dynamic hedging, you typically have moderate cost, although it really depends who is doing the hedging. I've seen examples of direct companies where dynamic hedging can be more expensive than any of the other strategies, depending on what view people have taken and how wrong it has become given the lack in discipline and improper corporate governance. Properly done, the over- or under-hedging is actually dynamically managed so that it does not get out of control.

  Furthermore, counter-party risks associated with this strategy are kept to a minimum since the trades put on are short-dated using very liquid instruments, which many times trade on the exchange.

In an ideal scenario, if you do not see any static solutions or reinsurance at a price that you want, the approach would be to first embark on a dynamic hedging program. This program would continue to be implemented until you see static solutions and/or reinsurance becoming available at a relatively attractive price—in which case, you would appropriately unwind your liquid positions and purchase these long-term assets. On the other hand, should these static solutions or reinsurance not become available during the life of the liabilities, one would continue with the dynamic hedging program until the expiry of the liabilities.

To illustrate, let us look at a numerical illustration associated with an annual ratchet product. Assuming that

- Policyholder is 60 years old
- Lapse rate is 5 percent

the following would be the costs associated with each strategy:

- Naked - Potentially 0 bps (with severe exposure)
- Reinsurance - Approx 30 bps
- Static Hedging - Approx 40 to 50 basis points (as it really depends on who's making the markets and what kind of a risk that they have on their existing books). The reason for this stems from the fact that investment banks may be already loaded up on one type of risk on their books. Therefore, to dissuade you from giving them more of that business, investment banks would slap in a premium to take on more of the same risks. On the other hand, it is possible to see investment bankers sometimes coming in at attractive prices. The reason for this stems from the fact that in doing a trade with you, the risk acquired would help offset some of the

existing risks they have on their books. As such, the hedger is actually paying the investment banks to get rid of their risks.

- Dynamic Hedging - Approx 18 basis points (if you know what you are doing). While it is quite easy to see the pickup associated with moving from a static to a dynamic strategy, it is important to note the other sorts of risks associated with it.

The risks underpinning all equity-based annuity products can easily be broken down into four categories. These can be classified generically as follows:

- Market - This relates to the market risk (i.e. how markets move and the impact of such movements).
- Actuarial - This relates to policyholder behavior and mortality.
- Regulatory - This relates to all risks underlying regulatory, accounting and legal issues.
- Operational - This relates to all risks underlying operational issues. It is one of the biggest understated risks since people are usually quite focused on market risk as it pertains to the hedging. Based on my audit experiences, I have found a lot of operational risks inherent in a hedging program. Examples of these are:
    - Key man risks: There's only one person in the company doing it with no trained backup.
    - Lack of documentation: The risk-management process is not being documented (i.e. nobody else in the company knows what that person is doing).
    - Lack of proper systems and controls: While trading off spreadsheets may be cost effective, the problem is that the person who is actually executing the trades is also the same person who is programming those numbers and reporting the risk exposures to senior management. As such, the thief and the police is the one and same person! Although most people are honest, we should not assume this to be case given the frauds and blow-ups happening quite regularly in the financial markets.
    - Data Quality: Any risk management program can only be as good as the quality and the timeliness of data used. In my experience, I have witnessed variable annuity writers of ratchet products not knowing if a ratchet has actually happened for each policyholder! I have also seen situations where brokers decide to do a favor for their clients by re-setting the timing of the deposits (or withdrawals) so as to ensure that their clients obtained a better return on the investment.
    - Controls/Governance: Trading limits in any active risk management program is very crucial. An example of this is the delta limit where one ensures that that fluctuations between the asset/liability delta stays within a pre-specified dollar value. So the person executing the trades knows a priori how much risk or market view can be taken. Similar philosophy holds for limits on gammas, lapses, volatilities, etc. These notional limits can also be easily extended and applied to value-at-risk,

earnings-at-risk, profits-at-risk metrics—depending on how much of risks one wants to stomach. An innovative method of introducing some of these metrics involve determining a priori the value-at-risk number, for example, and then deciding what kind of a likelihood you want assign to this risk number being exceeded (i.e. what is the probability of exceeding this limit). As an example, the corporate governance would dictate that with only a 5 percent probability you want to be able to exceed the value of risk limit of $10 million. When managing your risks on a day-to-day basis, as soon as this probability exceeds 5 percent, the risk manager would implement trades so as to bring the probability down to less than 5 percent.

- Consistency Between Hedging and Reporting: It is important to be able to get some consistency between what is being reported and what is being hedged. As otherwise, although one's risks may be economically hedged, the reporting of the financial results may depict otherwise! Given the current state of affairs pertaining to the treatment of hedges and derivatives for variable annuities, it is quite difficult to economically hedge the risks and be able to report exactly the current risk position—especially when derivatives are involved. As such, while you can be optimally hedged in one sense, it introduces noise to other accounting and regulatory metrics. So the trick is to decide on one's objective and hedging program, following which one needs to understand the ramifications of that implementation. I had the fortune of seeing firsthand the implementation of hedging programs where risks were economically hedged with the CFO freaking out a year later because the reported earnings continued to be volatile —not realizing that hedging economic risks does not imply smoother reported earnings!
- Benchmarking: While hedging is important, it is just as important to be able to benchmark the effectiveness of the hedging strategy, which only a few people are doing. The reason for this stems from the fact that you do not want to waste your money and resources, and if the implemented hedge strategy is not working for whatever reason, you want to know. As such, one needs to find an objective way to quantify how much risk has been taken, if the risk manager is doing a great job and whether as a firm you're getting compensated for the risk that you've taken.
- Attribution Analysis: When hedging, it is also important to be able to do good attribution analysis. The reason for this arises from the fact that no hedging program (except complete reinsurance) is going to be perfect. As such, one would need to be able to explain the factors driving the slippages in the hedging program, be they lapse assumptions, basis risks, poor risk management, etc. More succinctly put, you would need to identify if the slippages in your hedging program are due to the poor underlying assumptions or poor risk management.

Unlike insurance companies, banks spend quite a bit of effort trying to minimize the operational risks and keep them controlled. In many of the insurance companies running a hedging program, I have found that there is no emphasis on operational risks as people tend to forget that everything can just blow up even if your hedging program is good (as evidenced by the publicized disasters). The key to managing operational risk is to educate your senior management and board on exactly the types of risk you're managing and how you're doing it. As I said earlier, the unfortunate thing is that you have 15 minutes in a board meeting to talk about everything, and naturally it's not enough. But the flip side of it is that you need to get your board onside before you can actually do something. As a result, this becomes a balancing act of trying to educate the board members without overwhelming them with details.

Sometimes, I do get asked as to what a hedge report should include. As shown in Figure 18, it should summarize at a high level what kind of risk you have on your books, what is the mark-to-market value of the risks, what kinds of instruments are being used, what are the assumptions underlying the mark-to-market value, etc.

Figure 18

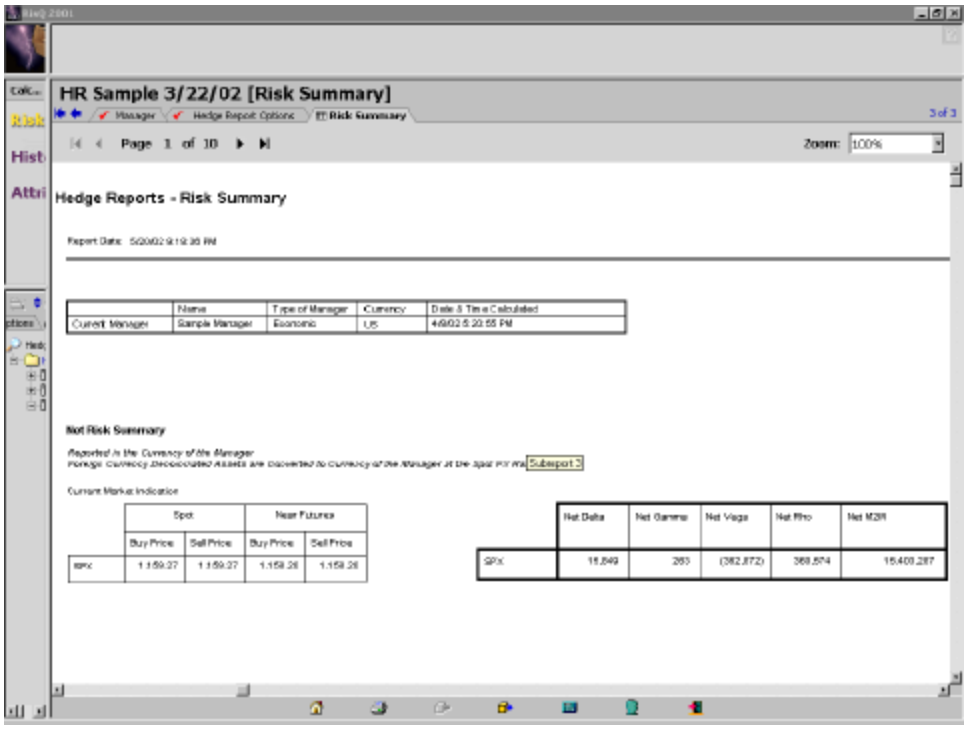

When managing market risk, one has to think about model risks and market gapping risks. Running a dynamic hedging program involves hedging your delta, gamma risks (just to name a few). But these are risk measures that are calculated for an instantaneous market change. In reality, markets do gap. As such, in implementing the hedging program, one has to consider such gapping risks. To do this, one has to look at a wider scenario analysis (e.g. markets running up by 5 percent, 10 percent, 20 percent) and then quantify how much risk you are exposed to. In addition to this, you also want to be able to quantify the probabilities associated with these scenarios. With these metrics, one would have enough view to know the likelihood of a scenario happening and the exposure under such a scenario.

When hedging, one is inevitably exposed to risk of being able to time your trade executions. In running a hedging program, contrary to what I have seen done in many insurance companies, you CANNOT have a rigid program where trade executions can only happen every Monday, Wednesday or even the third Wednesday of the month. Markets do move constantly. As such, any time that you can take advantage of favorable market conditions will result in a cheaper trade execution. Otherwise, it is possible that on the day of your trade execution, markets may have moved unfavorably making the execution an expensive one. So trying to dynamically understand markets and execute trades under favorable market conditions can put one ahead in terms of the amount of money saved. In

addition to this, being able to take market views also becomes quite important in terms of cheapening the hedge costs.

Finally, it is important to be able to quantify the risks associated with policyholder behavior. The more one can quantify the risks, the better handle one can have on risk management program—especially since these risks cannot be easily hedged!

MR. GILBERT: That concludes our presentation. We have some time for questions.

FROM THE FLOOR: Regarding the natural offset between equity-indexed annuities and variable annuity products, have you seen companies going naked on their equity indexed annuity to offset part of their variable annuity risk?

MR. GILBERT: I haven't seen it explicitly. I have seen some companies doing a very good job at hedging their equity-indexed annuity exposure, but then are not hedging their variable annuity block of business. In fact, by virtue of the fact that they're doing such a good job hedging the equity-indexed annuities, they are actually increasing their total company risk, because they're eliminating some of the offsetting exposures that were beneficial. But I haven't seen companies that have been measuring that and doing that rigorously.

MR. ZIMMERMAN: There's one cautionary principle I'd like to throw out, because that was something we considered as we wrote equity-indexed annuities. The variable annuity tends to offset the equity-indexed annuity Delta. It flattens your Delta, so you need fewer futures. But it doubles your Vega. So if you're just using one to offset the Delta of another, and the end of the quarter comes along and there's a spike in volatility to 50 percent, as you value your liabilities you're going to have a huge loss that quarter. You're just managing one of the Greeks and not all of them. Yes, it is a natural hedge for Delta, but there's also an interest rate risk, and volatility risk, and Gamma risk. So if that is something you're considering, you need to take some of Ravi's comments into consideration about your risk limits for volatility and interest rate and the other Greeks.

FROM THE FLOOR: So your company hedges each of them separately?

MR. ZIMMERMAN: We decided not to get into variable annuities, so we just have the equity-indexed portfolio.

MR. RAVINDRAN: If I may also add one more comment, the Deltas on the equity-indexed annuities tend to be unidirectional. You will find that the Deltas on variable annuities don't happen to be that way and it depends on the types of products that you have. I have seen cases of products that act like puts one moment and like calls the next. As such, in combing all equity-indexed annuity and variable annuity products, one needs to factor all these risk characteristics carefully.

FROM THE FLOOR: The question is for anybody on the panel. Ravi talked a little bit about operational risk, and what the banks do for operational risk, and I think the business put out a paper on operational risk. I just wondered if there's anything analogous to variable annuities and equity-indexed annuities as far as looking at operational risk?

MR. RAVINDRAN: I'm glad that you asked the question, because I am putting together a paper to cover the operational risks associated with hedging variable annuities and equity-indexed annuities. If you're interested, please give me your cards and I'll send them off to you. But to the best of my knowledge, I do not know of any such literature that is present.

MR. ZIMMERMAN: To share some more of Americo's mistakes with you, when we started writing this we were producing the sales report once a week. We didn't have any rollover, so we didn't think about having to quantify that. We didn't get any reports on surrenders or deaths or anything like that. Then we found out that we had a lag in reporting. We had the applications waiting to get on the system for two or three weeks, sometimes up to six months, which was shocking to us. We had to rattle some cages and let people know that was completely unacceptable. As we evolved, as we got to managing the risk through a dynamic hedge portfolio, we went to our IT department and asked for the cost to get a daily balance sheet. It was $50,000 annually, and missing one 10-point move in a day was $100,000 for us. So it was a no-brainer; the cost was justified. We had to do it.

FROM THE FLOOR: A quick question for Charles on the risk-neutral versus-real world scenario question. Are the risk-neutral scenarios really a practical option? If the market is in short-dated options, and we're talking about fairly long-dated guarantees, can you really construct risk-neutral scenarios that match our products? Aren't you forced to put some real-world views in there to extend the time period?

MR. GILBERT: That's a good question. We talk about manufacturing cost and whether or not you're able to replicate the payoffs of these long-dated options. That may not be possible in the capital markets. It's very difficult to find counter-parties, for example, for a 10-year put, let alone going out longer than that. But from a theoretical standpoint, the risk-neutral valuation would still be relevant in terms of giving you the theoretical capital market cost based on the capital market assumptions and assuming a risk-free rate. It's still worthwhile. Whether or not you could actually go out and buy that is another issue.

FROM THE FLOOR: But even if you can get one bid on a long-dated option, is that really a market price?

MR. GILBERT: Ravi, do you want to answer that?

MR. RAVINDRAN: I guess the short answer would be yes, but the long answer would be no. The reason why the short answer is yes is due to the fact that at the end of the day when someone comes to you and wants you to unwind everything as of today, given that there may be only one market maker, it is at that price that you have to unwind your positions. I am currently having a major debate with an investment dealer who is making markets in long-term equity options. They felt that in trying to take a view in markets one should not reflect that view when calculating mark-to-market values. For example, if one thought that the long-term volatility is not 30 percent but rather 18 percent, then that's a view, and therefore one would have to capture that view relative to where the capital market is. But trying to embed your view when marking-to-market your positions would be actually quite difficult to do in practice. If I had purchased some options two years ago when there were a lot more people making the markets, and all of a sudden I want to liquidate those options, who is there to buy those options up? If there's going to be only one person, then that becomes your market to deal with. It's a very nebulous situation, but having said that, you have to do your mark to market relative to where the implied markets are, which you can get by calling some market makers for long-term volatilities. But then, to incorporate your view in a real world sense, (e.g. volatility is only 15 percent) should be done via your trading limits as opposed to the way you mark to market or marking to market using a 15 percent volatility.

FROM THE FLOOR: If you're saying there's only one person in market, that's one person's view of the market. You're another person, if you have a different view and you choose not to sell it, then you don't have a price at all.

MR. RAVINDRAN: Exactly. But the key difference between you and them is that they are the "market makers," while you are just a hedger. That's the driver here. Because you're absolutely right, you can argue by saying I don't have to execute this trade. That's perfectly legitimate. So that's why I said you want to mark to market relative to your implied market conditions but incorporate your view using your trading limits. As such, if you have not breached your limits, there is no hedging to do.

MR. GILBERT: Just a quick poll, a show of hands to see what risk management strategy you're currently using. First of all, how many variable annuity writers are out there? Almost all the audience. Out of those that raised your hands, how many of you are currently running the risk naked? Okay, a small fraction. So the rest of the audience is going to raise their hands for reinsurance. How many are using reinsurance on new business? A smaller fraction still. How many are using some form of hedging, either using a static or dynamic hedging approach? Three or four people. Okay, very good.