RECORD, Volume 27, No. 3*

New Orleans Annual Meeting
October 21-24, 2001

Session 160TS
Understanding And Managing The Risks Underlying Guaranteed Benefits In Variable Annuities

Track: Product Development

Moderator: INGER S. HARRINGTON
Instructors: DAVID L. BRAUN
CHARLES L. GILBERT
INGER S. HARRINGTON

Summary: In addition to providing the opportunity to allocate deposits among various investment categories, variable annuities now include benefits that limit the investment risk under specified investment return circumstances. Guaranteed minimum income benefits (GMIBs), guaranteed minimum accumulation benefits (GMABs), guaranteed minimum death benefits (GMDBs), and death benefits designed to cover the approximate income tax on the gain in the contract all provide additional protection for the purchaser, but they may also bring added risk to the insurer. The measurement and management of the underlying risk factors is the focus of this session.

MS. INGER S. HARRINGTON: There are really three things we're going to be talking about. Charles Gilbert will focus on dynamic hedging. David Braun will focus on natural hedging, which is something that hasn't been discussed in great detail at the conference. And then I'll be talking about reinsurance—how to find reinsurance and what approaches that have worked very well for our clients.

Let me just tell you a little bit about our speakers. Charles Gilbert is part of the financial services practices of Tillinghast-Towers Perrin in Toronto. He's active in risk management topic areas. He's also the leader of his firm's equity risk management initiative for North America.

Note: The chart(s) referred to in the text can be found at the end of the manuscript.

David Braun is with the Hartford office of Ernst & Young. During David's three-plus years at Ernst & Young, the majority of his work has focused on risk management and stochastic modeling for life insurance and annuity companies. Prior to joining Ernst & Young, David spent nearly six years at Cigna, where he held various positions in pricing and financial reporting.

I'm Inger Harrington, and I work for a company called Ruark Insurance Advisors. This is a company that does both consulting and reinsurance intermediary work for the variable-annuity market. Prior to joining that company, I worked with Cigna. My last six or seven years have been focused on the reinsurance of variable annuities.

MR. CHARLES L. GILBERT: I want to review the risks associated with variable annuities, describe the risk management process that's typically used, and then review the different risk management options. The focus of my presentation is dynamic hedging. I'll provide an overview of the dynamic hedging process, walk through a numerical example, and then briefly touch on some of the practical issues associated with dynamic hedging.

In terms of the risks associated with variable annuities, basically what we have here is a product with investment guarantees that is effectively an embedded derivative. The insurer has written a put option. The other thing is that the premiums, the mortality and expense (M&E) charges, are based on the fund size or assets under management. Variable annuities can have product features that may be quite complex and difficult to understand. So there are a number of risks associated with these types of products.

Earnings volatility can result from either real economic losses—the payout under the guarantees—or the revenue loss from the fee income collected on lower fund balances in declining markets. Earnings volatility can also result from artificial accounting treatment. This would happen when you have one side of the balance sheet being marked to market, the other side being fixed at book. Even though you don't have a real economic exposure, you're still seeing some of this volatility in the emergence of earnings, in the emergence of accounting earnings.

There's also regulatory risk. This is something we've seen in Canada, where the Office of the Superintendent of Financial Institutions (OSFI) has recently implemented capital requirements for segregated fund products that effectively have a guaranteed minimum accumulation benefit. This was not necessarily anticipated at the time that insurance companies had priced their segregated fund products. As a result, a lot of them find that without having priced in the cost of capital, they have products that aren't as profitable as they once were.

The capital requirements can be significant. It's going to be a function of the guarantee, the level of the guarantee, the guarantee features, the in-the-moneyness of the guarantees, where the market is, and how long the product has been on the market, the duration of the business. In Canada the average is probably somewhere between 5 and 8 percent of the assets under management. I think that was at last year-end, and this year-end is probably going to be higher, given the markets now. In addition, companies that have to put up that level of capital are also faced with another issue, which is capital volatility. So from year to year, depending on where the market is and again where the capital—or where the in-the-moneyness of the guarantees are—that's going to play on the level of capital that'll be required.

Other risks include operational risks and legal risks. A number of companies have offered variable annuities with enriched guarantee features, resets, ratchets, and other features that the administration systems aren't necessarily capable of administering. We've seen one company in particular that offered resets on its product, but its administration system couldn't track the resets. Also, the product is not always well understood by the policyholder. So there are some issues with respect to making sure the product is sold properly.

To give an example, Chart 1 shows the risk profile of a guaranteed minimum accumulation benefit (GMAB). The red line shows the unhedged risk profile. This is for a $1 billion in-force block, and the red line shows the distribution of profits and losses. So, if you look at the extreme left tail exposure, the present value of the deviation from expected on this block was $175 million, negative. In comparison, the hedged position—this would be using dynamic hedging—is the blue line. You see that you haven't totally eliminated the left tail exposure, but you significantly mitigated that. In comparison to the $175 million loss, you have a $25 million loss in the dynamic hedging example.

The risk management process involves five basic steps. The first step is understanding the risk and quantifying the risk exposure, looking not only at the expected cost but also at the risk profile and understanding the tail exposure that you're facing. The second step is deciding whether that risk exposure is appropriate. An important point to note is that life insurance companies and financial institutions in general are in the business of taking risks, of writing risk, and the goal of risk management is not necessarily to eliminate or even minimize your risk. It's going to be a function of the company's risk tolerances and constraints. The third step is analyzing all of the various risk management options, looking at both the cost and the effectiveness of those risk management options. The fourth step is formulating and implementing your risk management strategy and then, once you've implemented that on an ongoing basis, monitoring the risk exposure and the results of your risk management strategy.

Just to recap what some of the risk management options are, we've broken this down into two basic categories –(1) risk management strategies in which a passive approach is taken and (2) risk management in which an active approach is taken (Chart 2). The passive approach doesn't necessarily mean that the company is doing nothing. Running the risk naked can be a legitimate risk management strategy. If the company is monitoring its exposure, quantifying its exposure, and if it’s within its risk tolerances and constraints, then running the risk can be a perfectly legitimate risk management strategy. You can run the risk naked, or you can run the risk and hold capital, although, as I mentioned, in Canada, if you’re going to run the risk naked, you have to hold capital, and I think that’ll soon be the case in the U.S. as well.

I won’t go into too much detail because we’ll talk about these later, but more active approaches would include looking at the overall distribution of the product offerings; the specific product design features—this is something that David Braun will talk about a little bit later, specifically how you can identify natural hedges and offsetting exposures that would exist in your portfolio; reinsurance, which Inger will discuss in more detail; and then hedging strategies, both static and dynamic.

Now I’d like to discuss some of the advantages and disadvantages of each of these strategies. Among the advantages of running the risk naked, which is often the do-nothing or default approach, is that it’s easy to implement. There’s no upfront cost associated with reinsurance premiums or hedging costs, and running the risk will have the highest level of profits on an expected basis, ignoring the cost of capital.

The disadvantages of running the risk naked are that you potentially could have significant earnings volatility. There’s exposure to large and potentially catastrophic losses. Company management may be perfectly comfortable with this strategy if you have the highest expected profits. But if there’s a 1 percent chance that you’ll blow up the company, maybe you’ll want to reconsider that strategy. Now more than ever, running the risk naked will attract increased attention from regulators and rating agencies. We’re seeing that on both sides of the border, and there are, as we’ve discussed, potentially high capital requirements.

The advantages of reinsurance are that it’s easy to obtain internal approval. It’s something that’s familiar. It’s something that’s easy to implement. It can be customizable. Both the insurance and the financial market risks are covered, and it provides some certainty of price for new business, unlike hedging or dynamic hedging and reinsurance, where the cost is not known at issue but only at the end of the expiry of the contract.

The disadvantage of reinsurance is that it’s expensive. It may hurt your competitiveness. There’s limited coverage, especially for the tail exposure, and this is an important point: a number of companies are really interested in reinsuring their tail exposure. They may be comfortable with the first layer of reinsurance, but that’s not what the reinsurer’s prepared to offer. So what the company is often getting is that it’s giving up its profits. It’s paying for reinsurance premiums on coverage or the risk that it would be comfortable accepting and not having the coverage on the part that it really wants to reinsure.

There is counterparty credit exposure with reinsurance. Companies have no control over the reinsurer’s risk management, and unlike mortality risk, the nature of equity market risk is such that the traditional concept of pooling does not apply. If you aggregate this type of exposure, you have a concentration of the risk as opposed to diversification of the risk. There's an argument to be made that you can diversify some of the basis risk, but the equity risk is increased. Reinsurance is irreversible. It's illiquid. And as a number of companies have learned first- hand, the renewal of the coverage is not guaranteed.

I'll talk a little bit about static hedging and dynamic hedging. Static hedging typically refers to strategies involving long-term, over-the-counter (OTC) options in which there's little or no rebalancing, whereas with dynamic hedging, there's a rebalancing of the position on a more frequent basis. The advantages of static hedging are that there's little or no ongoing rebalancing. There are limited internal controls because this is a buy-and-hold type strategy. And there's limited exposure to actual volatility. You're basically paying a premium for the implied volatility, but in exchange, you're not exposed to actual or realized volatility. One of the disadvantages of static hedging is that it's expensive. There's a lumpy premium flow. What we mean by that is that typically there's a critical mass or a minimum contract size that you need in order to buy the OTC options to structure them. So, between the time that you receive the premiums and the time that you've accumulated enough critical mass, you have some residual exposure that you either carry naked or you find a way to hedge in the interim.

There are a limited number of providers, especially at the longer terms, and there's significant counterparty credit exposure. A number of insurers won't deal with anyone who is less than AAA rated. There's limited liquidity. There are no established secondary markets for these types of contracts. So, if you need to unwind an OTC contract, generally the only party that you can find to do that will be the original counterparty. Static hedging cannot address variances in expected persistency. If you're using a rollover strategy, there is some uncertainty of the future prices and the availability for those rolled positions. You are exposed to increases in implied volatility, and there's also basis risk.

The advantages of dynamic hedging include that it easily addresses variances in your lapses and market conditions. The ultimate cost is a function of actual volatility, which is significantly cheaper than the implied volatility that you would pay if you were statically hedging. It uses the most liquid hedging instruments that are generally exchange traded, and as a result, there's limited counterparty credit exposure. Again, as David will talk about, there are opportunities for product design initiatives in which you can identify offsetting option sensitivities and create offsetting liability exposures.

One of the disadvantages of dynamic hedging is that it's complex. It's difficult to explain and understand. There are a lot of internal approvals required to implement a dynamic hedging program. There are ongoing controls and reporting that will be required. There can be some operational risks of trading execution if you don't have the proper controls in place. The ultimate cost of dynamic hedging is not going to be known until the expiry of the contract. In general the expected cost of dynamic hedging will be lower than reinsurance or static hedging, but the ultimate cost could be higher or lower than expected. There will be some residual risks associated with dynamic hedging, especially if you're trying to dynamically hedge an actively managed fund. There'll be some basis risk, and sophisticated systems and expertise are needed to execute a dynamic hedging program.

I'd like to talk a little bit more specifically about dynamic hedging. So, what are we trying to do here? The objective of dynamic hedging is to replicate the option that has been sold in the liabilities. In the case of any product that has an investment guarantee, the insurer has basically written an option or sold a derivative that is embedded in the liabilities. Then, the goal is to buy an option or create a hedge portfolio that replicates that option and whose value is going to increase or decrease in the opposite direction for given changes in market variables. A successful hedging program will look at the hedge effectiveness and the hedging costs. One of the biggest challenges for companies that have implemented dynamic hedging programs has been performing accurate attribution analysis, to explain the source of gains and losses.

Table 1

Hedge Effectiveness

- Effectiveness = How close is (change in liabilities - change in assets) to zero
- Ineffective Hedge

Time Period	0	1	2	3
Liabilities (Embedded Options)	500	525	495	525
Assets (Hedge Portfolio)	100	110	105	140
Net	400	415	390	385

- Effective Hedge

Time Period	0	1	2	3
Liabilities (Embedded Options)	500	525	495	525
Assets (Hedge Portfolio)	100	122	93	128
Net	400	403	402	397

All amounts above are the MV of the assets and liabilities

14

We define hedge effectiveness as, how close to zero is the difference between the change in assets and the change in liabilities? Here's an example of an ineffective hedge (Table 1). You look –across time periods—and this could be days, weeks, or minutes; it doesn't matter—and you look at the embedded option in the liabilities.

We start with a value of 500. We call this the marked-to-market value of the liabilities. The marked to market increases in time period one to 525, –then decreases to 495, and then goes back up to 525 in time period three. In this example we've set up a hedge portfolio that tries to offset those changes in the liabilities.

The first thing you notice here is that the marked to market of the hedge portfolio doesn't equal the marked to market of the liabilities, and it doesn't need to. Here you're trying to create a hedge portfolio for which the changes are going to offset the changes in the liabilities, so that's what you're concerned about. The net isn't really relevant. We're not really concerned about the difference between the marked to market or the liabilities and the assets. We just want to track through time how close the change in the two has been. In time period one, when the liabilities increased to 525, the assets only increase to 110. The liabilities increase by 25, but the assets only increase by 10, so you have a loss of 15.

In time period two, the liabilities drop by 30. The assets only drop by five. Now you have a gain of 25. And in time period three, the liabilities increase by 30, but the assets increase by 35. So now you have a further gain of five. In this example you've actually generated gains on your hedging strategy. Cumulatively, you have a gain of 15, but that's not the point. The goal is not to try to generate gains and losses. We're not taking an active management approach to our hedging strategy. The goal is to get that as close to zero as possible.

In the other example, the effective hedge has the same numbers for the liabilities. The change in the assets, the hedge portfolio, is a lot tighter, and you have a net gain of three compared to the gain of 15. So we say that in the second example, you have a more effective hedge.

Just to go through the dynamic hedging process quickly, the first step is to identify and determine the hedge instruments that you're going to use in your hedging strategy (Chart 3). You're basically going to look at your market information and at the liquidity of the various instruments, the open interest, the correlation to the underlying that you're trying to hedge, and you're going to come up with which hedge instruments you're going to consider in your asset universe. The next step is the portfolio replication. You start off by looking at the funds that are available that are being offered to the policyholders.

One insurer that we worked with had more than 80 funds, and we wanted to group those to a smaller, more manageable number. We looked at grouping similar funds and then did the portfolio replication on a smaller subset. The next step is partitioning the liabilities into hedge baskets. –Again, you're going to start with your seriatim liability file. If runtime considerations are an issue, then you might want to consider grouping those liabilities. Once you've done that, then you partition those liabilities into the hedge baskets, which will be the indexes or the instruments that you identified in the first step. Once you've done that, then you can model the embedded liability options for each of those hedge baskets and then calculate the option sensitivities or the Greeks, and the Greeks are calculated on a real-time basis with a live feed from the markets.

You do the same thing for the assets. You look at the asset universe. These are the hedge instruments that you've identified that you're going to consider, and you calculate the Greeks for those assets, and now you're able to execute your hedging strategy. Before you can execute your hedging strategy, however, you have to come up with your hedging strategy. In formulating your hedging strategy, you should do a couple of things. You should test different strategies. You should look at setting up a mock portfolio, testing on a live basis each of the strategies that you want to look at, doing some back-testing on each of those strategies based on historical market data, and then you should do what-if scenario analysis.

How well would that strategy have held up under that scenario? For each of your strategies, you want to evaluate the hedge effectiveness, the hedging costs, and you should perform your attribution analysis. Once you've selected and executed your hedging strategy, the result is your rebalanced hedge portfolio.

Now I'll talk a little bit more about each of these steps. In terms of determining the hedges, for instance, what are some of the questions that you need to ask yourself? The first question is, what is the underlying that you're trying to hedge? Then, how widely traded are the instruments that you're looking at? What's the open interest? Is there a danger that if this contract is too thinly traded that you're going to move the market if you use this instrument? How liquid are the instruments, i.e., how large is the bid-ask spread? How much basis risk exists if you choose an instrument that's maybe more liquid, for example, but not a perfect correlation to the underlying? How well does the hedge instrument replicate the underlying? Are transaction costs significant? Are there large, potential cash-flow requirements? Then, what is the hedging strategy? What option sensitivity are you trying to hedge?

Portfolio replication involves replicating the fund on which the guarantee is based, Basically this is a fund decomposition. For an actively managed fund you need to understand the investment policy for each fund. You may actually want to interview the portfolio manager with respect to how the fund is being managed in practice compared to what's stated in the investment policy. Generally there'll be ranges for the asset mix. –Historically, what has the style been for this manager? Then you want to do some statistical analysis of the portfolio replication based on historical data and recognize that your replication may or may not hold up over time. You want to track any basis risk that you have as a result of your portfolio replication.

To give a simple example, just look at two funds, Fund A and Fund B, both actively managed funds. After going through this process, you've identified that Fund A is effectively equivalent to 50 percent Standard & Poor's (S&P), 50 percent NASDAQ. That's the decomposition that you've come up with. As for Fund B, you've decomposed it into 25 percent S&P 500 and 75 percent NASDAQ. This is what you're going to be hedging. You're not going to be hedging the actual fund. You're going to be hedging the replicated fund. To the extent that you've got any basis risk, that's unhedged.

The next step is partitioning the liabilities into the hedge baskets. The first step is to perform any grouping of the seriatim file. We've looked at policies or in-force blocks that have had more than 300,000 policies, and we've been able to run that overnight. But if it's going to take you several days to model your liabilities, you may want to look at grouping your seriatim file. The hedge baskets correspond to the index or hedging instruments that you've chosen. Once the portfolios have been replicated using the hedge instruments, liabilities are partitioned into the corresponding hedge baskets, and the embedded options are modeled for each hedge basket.

Table 2

Partition the Liabilities into Hedge Baskets

- Perform any grouping of seriatim file (for run time considerations)
- Hedge baskets correspond to index/hedging instruments chosen
- Once portfolios have been replicated using hedge instruments, liabilities are partitioned into corresponding hedge baskets and the embedded options are modeled for each hedge basket

Example: Partitioning of Liabilities

Policy	Amount	Fund A	Fund B	S&P500	NASDAQ
1	1,000	10	90	275	725
2	2,000	50	50	750	1,250
Total	3,000			1,025	1,975

18

Here's a very simple example of partitioning the liabilities (Table 2) involving an in force of two policies. Policyholder 1 has invested an amount of 1,000; 10 percent of that is invested in Fund A, and 90 percent is invested in Fund B. We can come up with the amount of S&P and NASDAQ exposure for that first policy. Similarly, we do the same thing for Policyholder 2, and we come up with the partitioning of the liabilities into the two hedge baskets. In this case, the hedge baskets are the S&P 500 index and the NASDAQ index.

In the next step, for each hedge basket, you need to model the embedded liability options. What you model is what you're going to hedge. Here, we're computing the value of the option and its sensitivity, The key question is what you want to hedge. Is it the economic value of the embedded derivative, or is it something else? Investment guarantees and certain product features such as resets involve a complex optionality for which a closed-form analytic solution may not exist. In those cases, you may want to use a Monte Carlo technique to calculate the value of the option and its resulting sensitivities or Greeks.

So, what are the Greeks? The Greeks measure the sensitivity of the option to changes in key market variables. Delta is the partial derivative of the option value with respect to the underlying asset price or, in plain English, it measures the change in the price of the option to changes in the underlying. Gamma is the second partial derivative of the option value with respect to the price of the underlying asset. Vega measures the change in the price of the option to the change in volatility. Theta measures the change in the price of the option with respect to the decrease in time to maturity or time decay. Rho is the change in the option for changes in the risk-free rate.

Table 3

Theory - The Greeks

$$\Delta\Pi = (Delta)\Delta I + \frac{1}{2}(Gamma)\Delta I^2 + (Vega)\Delta\sigma + (Rho)\Delta r + (Lambda)\Delta L + \ldots$$

$\Delta\Pi$ = change in Liability
ΔI = change in underlying
$\Delta\sigma$ = change in volatility
Δr = change in yield
ΔL = change in lapses (if appl)

Equities	Fixed Income
Delta	Duration
Gamma	Convexity

- similar A/L approach to match instantaneous price sensitivities

- When the # of immunized Greeks increases
  - volatility in A/L portfolio decreases
  - instruments used increases
  - hedging program is said to be more "static", less "dynamic"
- Delta hedging is the simplest form of dynamic hedging

21

If you look at the option sensitivities, if you look at the Greeks, you can use a Taylor-series expansion to come up with your approximation for the change in the embedded option that you're valuing (Table 3). In this example, the change in the embedded option is going to equal delta times a change in the underlying, plus one-half gamma times the change in the underlying index squared, plus vega times a change in volatility, plus rho times the change in the risk-free rate, plus—we've defined this other variable lambda times the changes in lapses or other actuarial assumptions. If you want an analogy, you can look at hedging equity risk and look at delta and gamma as being equivalent to hedging interest rate risk, where delta would be the analog of duration and gamma would be the analog of convexity. The more you want to immunize, the more Greeks you want to hedge, the greater the number of instruments you're going to need but the less volatile the asset/liability portfolio will be.

The next step is determining the hedging strategy. This depends on a number of factors, including the sophistication of the company, the expertise that the company has; the risk tolerances, i.e., the desired hedge effectiveness; and the amount the company is willing to pay for the hedge. Dynamic hedging provides a company with a lot of flexibility in terms of how much risk mitigation you want and how much you're willing to pay for it. The hedging cost, as we mentioned, is going to be a function of the realized volatility, and it's not going to be known until the end of the maturity date of the contract. Obviously, you can't have both—there's a tradeoff between the hedge effectiveness and the hedging cost.

Table 4

Hedging Example - Hedging Delta, Gamma and Vega

- Step 1: Calculate Liability Sensitivities to the S&P500 index

Delta	Gamma	Vega
-46	0.2	195

- Step 2: Find S&P500 options to hedge both Gamma and Vega

Ticker	Instrument	Maturity Date	MV	Delta	Gamma	Vega
SPX 6 P995	Put Option	15-Jun-01	21	-0.197	0.0016	1.514
SPX 6 P1150	Put Option	15-Jun-01	75	-0.534	0.0026	2.149

* Source: Bloomberg Professional Service as of 4pm, March 22, 2001

- 195 = X*Vega (SPX 6 P995) + Y*Vega (SPX 6 P1150)
- 0.2 = X*Gamma (SPX 6 P995) + Y*Gamma (SPX 6 P1150)
- Solve for X and Y
  - X = 155 and Y = - 18
- This is usually an iterative process since hundreds of options are available.

23

I'd like to go into a numerical example (Table 4). This looks at hedging three Greeks—delta, gamma, and vega. We'll start by looking at the liability sensitivity of a put option on the S&P 500 index. Here we've modeled the liabilities. We've calculated the Greeks. The Greeks are negative 46 for delta, 0.2 for gamma, and 195 for vega. The next step is to find the two options that are going to hedge your gamma and your vega exposure. We looked at the option chain. We selected two put options, one with a strike of 995 and one with a strike of 1,150. Both have the same maturity date. The delta/ gamma/vega and the market value of the options are shown in the table.

We want to get ourselves to a neutral gamma and vega position. In the first step, you're basically solving two equations and two unknowns. You want to know how many of the first put option you want to buy and how much of the second option you want to buy. If you go through the math, solve for x and y, you find that you want to buy 155 contracts of the first put option, and you want to sell 18 of the second contract.

Table 5

Hedging Example - Hedging Delta, Gamma and Vega

- Step 2 (Continued):
  - Residual sensitivity

Delta	Gamma	Vega
-25	-0.001	-0.988

- Step 3: Use futures to hedge remaining delta risk

Ticker	Instrument	Maturity Date	Delta	Gamma	Vega
SPM1	Futures	1-Jun-01	1	0	0

* Source: Bloomberg Professional Service as of 4pm, March 22, 2001

  - Residual Delta = -46 - [X * Delta (SPX 6 P995) + Y * Delta (SPX 6 P1150)]
    = -46 - [155*-0.197 + -18*-0.534] = -25.08
    = Short 25 futures contracts
  - Residual sensitivity

Delta	Gamma	Vega
-0.08	-0.001	-0.988

- Rebalance over time or after material changes in underlying, volatility or interest rates.

24

Now, you look at the residual sensitivity, and you see the gamma and vega exposure has been significantly reduced, almost to zero (Table 5). In the third step, you want to use futures to hedge the remaining delta exposure that you have. If you look at how we got that negative 25 delta, the calculation shown in Table 5, you had an original delta of 46, and then if you looked at the delta of the option contracts that you purchased, you get a residual delta of minus 25. What that's telling you is that you need to short 25 futures contracts, and the residual sensitivity is shown in Table 5. You see that you've been able to hedge your delta/gamma/vega exposure pretty close to zero. –Again, this is an ongoing process. As the name dynamic hedging implies, you want to rebalance this continuously to maintain that neutral position.

Earlier, we talked about static hedging versus dynamic hedging. One form of dynamic hedging that's quite common is delta hedging. Basically, delta hedging is a buy-high/sell-low strategy. The greater the volatility, the greater the price of the delta hedging strategy. When you buy a static hedge, you pay a premium. Basically you're paying the implied volatility. If you look at Chart 4, the blue line is the volatility index (VIX). The VIX Index basically measures the implied volatility of the S&P 100 Index on its 30-day implied volatility. And the red line is the actual or historical volatility, 30-day volatility.

If you look at the difference between those two curves, you see that the average spread is about 5.8 percent. So, 5.8 volatility (vol) points is pretty expensive if you're dealing in option premiums. One of the reasons for delta hedging would be to try to save on the difference between the implied and realized volatility. Union Bank of Switzerland (UBS) did a study on long-term volatility and the spread between long-term implied vol and long-term realized vol and found that the spread was closer to 7 percent and that in the long term, implied volatility was mean reverting. A number of companies have looked at that fact and tailored their hedging strategies around that.

Another key practical issue is the trading frequency versus the transaction cost (Chart 5). More frequent rebalancing will improve the hedge effectiveness, but it'll also increase the transaction costs. One of the things that is unique about these types of products with these types of guarantees is the long-term nature of the liabilities. When you're looking at your strategy, one of the things that you may want to consider is the long-term nature of your horizon and whether or not you want to consider this and test different strategies over different types of market conditions.

One example would be if you adopt a frequent trading strategy. Look at the events of September 11, when the markets dropped by 15 percent in one day. If you were delta hedging, you would have rebalanced and ridden the markets all the way down and then all the way back up. It would have been very expensive. If you had waited and rebalanced on a weekly or monthly basis, you wouldn't have incurred that cost because the markets would have come back up. Obviously there are some risks associated with that, but if you have a longer-term horizon, it may be something that you want to test.

What are some things of which you should be aware? We talked about the rebalancing frequency versus transaction costs, the volatility premium, and the static hedges. One thing that you need to be aware of is the trading behavior of the capital markets. There have been points in time when it has been impossible to trade because of what are known as liquidity holes, which are characterized by significant volatility in the markets. The bid-ask spreads widen tremendously or explode. You're unable to trade. Those are things that you have to take into consideration when you're determining your hedging strategy.

What if something like this happens? What's my exposure? You should test those scenarios. In the simple example that I discussed, when we looked at hedging vega, a simplifying assumption was made that basically volatility was constant and level. In reality, volatility has a surface. If you look at the term structure of volatility, you really need to be hedging across all of the points, not just making a simplifying assumption that your volatility is going to remain constant.

There are basis risks. There are data errors and assumption changes. If you implement a hedging strategy and then find that there are data errors in your liability file or that you need to change your actuarial assumptions, then that's going to affect the value of your embedded option. It's going to require you to rebalance your portfolio, and any resulting gains and losses are going to flow straight to the bottom line.

Policyholder behavior can be an important consideration. It's difficult to model. There's model risk. One of the biggest challenges that we're facing currently is a lack of understanding of the exposure to different trading strategies. I think there needs to be a greater awareness of what this particular risk management strategy can and cannot do for companies. There are the operational risks that we mentioned earlier.

So, how do you avoid the pitfalls? You need to use advanced tools to model these exposures and the trading strategies. You need to conduct proper testing of the strategies themselves. Test turbulent regimes and normal regimes. You need to do back-testing against historical data, what-if scenario analysis, and mock testing of hedging strategies in live markets. You need to identify and quantify unhedged basis risks and put in place effective controls, procedures, and reporting, including proper attribution analysis. Senior management and the board need to be aware of and educated on processes, strategies, and exposures. There are some pitfalls with dynamic hedging. There are some things that need to be carefully considered and addressed, but if you do these things, you can avoid these pitfalls, and dynamic hedging can be an effective risk management strategy for these types of risks.

MR. DAVID L. BRAUN: I'm going to be talking to you about the concept of using natural hedges to mitigate the risk embedded in today's variable annuity (VA) guaranteed benefits. The concept of a natural hedge should not be foreign to an insurance company. Insurance companies have historically made a lot of money by pooling risk exposures and profiting from the reduction of risk created by risk pooling/diversification (e.g. the premium from life insurance policyholders that do not die, help pay the claims of those that do die). Thanks to the revolutionary work done by Harry Markowitz on the topic of Modern Portfolio Theory, all investment professionals know about and exploit the power of risk diversification. Risk diversification can be thought of as combining two or more risks that have less-than-perfect correlation, or ideally negative correlation, in order to lower your overall risk exposure. My presentation is going to deal with this topic as it relates to the guaranteed benefits in VA products.

I thought I'd start with some general background information and then jump into a case study that will hopefully illustrate the point. There are two broad risks that VA writer faces. The first is revenue risk, which is basically the risk that when the markets are down, your mortality and expense (M&E) revenue is down (since M&E is a fixed rate applied to the asset base which is now reduced due to market underperformance) and thus your earnings are down. The second risk is claims risk. This is the risk associated with the guaranteed benefits that the VA writer provides to the policyholder. The VA writer has essentially sold the policyholder an option on the account value (either a put or a call option depending on the benefit design). When the markets move against the VA writer, these benefits become in-the-money and the VA writer may be forced to pay claims that exceed the policyholder's account value.

Most VA writers that I have worked with have grown accustomed to living with revenue risk (albeit unwillingly). They look at a distribution of the present value of that revenue stream under thousands of stochastic scenarios and see quite a dramatic disparity between the good scenarios and the down scenarios. However, when they explore the cost of hedging out that volatility and trying to flatten that distribution, they often conclude that it is too expensive. The VA writer would have to forego the "home runs" that you hit in a bull market in order to afford the protection required in a bear market. Most senior management teams that I have spoken with are reluctant to do that, as it is very hard to explain flat earnings in a bull market when all of your competitors are experiencing strong earnings growth.

However, VA writers are continually looking for ways to mitigate the claims risk. Whereas revenue risk is often thought of as earning less money but still earning money, claims risk is often thought of as losing money. VA writers have traditionally used either reinsurance, as Inger is going to talk about, or a derivatives-based hedging strategy, which Charles just spoke about. What I am proposing here is the third possible risk mitigation technique to deal with claims risk. This alternative focuses on using another liability to mitigate the risk generated from the liability for which you wish to reduce claim risk. Our goal is to identify two liabilities with offsetting risk profiles and then balance the mix of the sales for these products such that your aggregate claim risk is reduced.

For the purposes of this presentation, I'm going to do something that I do not normally like to do. I usually believe in looking at risk from a holistic view, either the whole line of business, the whole company, or at a minimum, the whole VA product. However, in order to illustrate the concept of natural hedges, I will carve out only the guaranteed benefit claims piece and will ignore all other profit/risk sources from the product line, company, or VA contract.

In the case study, we are going to look at two of the hot riders that are currently being offered by VA writers. One is the tax relief rider (TRR), or earnings enhancement benefit (EEB) as some folks call it. The TRR is designed to reduce the tax burden to the beneficiary upon death of the annuitant. The beneficiary of the VA proceeds must pay income tax on the gains in the contract. The gain is equal to the amount disbursed from the contract less the premium put into the contract. In order to compensate (or "make whole") the beneficiary, the TRR will gross up the death benefit payment for the estimated taxable gain.

A common TRR design is to gross up the death benefit for an amount equal to 40% of the taxable gain. For instance, let's assume a policyholder dies with an account value of $150,000 and the premium was $100,000. The death benefit without the TRR would be the account value of $150,000 and the beneficiary would be taxed on the $50,000 of gain ($150,000 less $100,000). However, the TRR would gross-up the death benefit by $20,000 (40% times $50,000) and pay the beneficiary $170,000. The beneficiary will now be taxed on the $70,000 excess over premium (i.e. the new gain) and, assuming an average tax rate of 29 percent, would end up netting the original $150,000 after taxes. The VA writer is on the hook for this $20,000 gross-up amount.

The second guarantee in the case study is a GMAB. This is one of the popular living benefits that is out there, particularly in Canada. I have picked a simple design just to illustrate the point. With this GMAB, the VA writer has guaranteed that on the seventh policy anniversary, the policyholder will have no less than 110 percent of the original premium deposit. For example, if you put $100,000 into a VA with a GMAB and on the seventh anniversary your account value was only $95,000, the insurer would kick in $15,000 to get you up to $110,000 (which is 110 percent of your premium).

Many VA writers are reluctant to offer these two types of benefits due to their concerns about the risk profile of such guarantees. I plan on showing how setting up a natural hedging program can reduce the risk exposure of these products and thus allow a risk adverse insurer to go to market with the products. Charles talked a lot about the dynamic hedging using equity-based derivatives. I'm going to talk about delta hedging the risk exposure for the TRR using the GMAB and visa versa.

Delta is the change in the option value with respect to the change in the underlying. The underlying in this case study is the account value. Delta is analogous to duration in the fixed-income world, which you will recall is the change in a bond's price given a change in interest rates. I will use stochastic simulation to calculate the delta for my two guaranteed benefits. I will generate 10,000 stochastic trials, run both products through them, and figure out the average present value of claims that the VA writer must pay. That will be our baseline value. Then I shock the account value up 1% and down 1% (just like in duration when you shock the yield curve up and down by a basis point). Then I run for both of these through the same 10,000 scenarios that generated the baseline value and determine the average present value of claims in the +1% and –1% runs. We now have three different values for the present value of the claims (for each of the two products), from which we can calculate the delta for both products.

At this point we should elaborate on the embedded options within the TRR and GMAB. The TRR is equivalent to a call option. The buyer of the call option is the policyholder and the seller is the VA writer. The buyer makes money when the markets go up and thus the seller of the call loses money when the markets go up. Thus, this product has a negative delta from the insurance company's perspective. It loses money when the markets go up, so the insurer's profit goes in the opposite direction of the market (i.e. negative delta). In order for an insurer to hedge this, it would have to buy delta. It would go long in futures contracts or buy call options.

The GMAB is exactly the opposite. It's equivalent to a put option. The buyer of the put, the policyholder, is protected from downside movement in the market. If the market goes down, he or she earns money on the GMAB option (via the kick-in to their account value in the seventh policy year). The seller of the put option, the insurer, loses money when the market drops. Therefore, it has a positive delta. Its value to the insurance company and the account value move in the same direction. To hedge this exposure, the insurer would have to sell delta by entering a short position on futures contracts or buying put options. Before we get into the delta hedging exercise, I would like to further illustrate how the risk dynamics of the TRR and GMAB are opposites.

Table 6

Year	Annual Return	Cumulative Return
1	27%	27%
2	60%	104%
3	44%	192%
4	20%	251%
5	24%	337%
6	32%	478%
7	-8%	429%
8	38%	632%
9	-3%	611%
10	11%	691%
11	13%	793%
12	-7%	727%
13	12%	826%
14	17%	979%
15	38%	1391%
16	15%	1609%
17	0%	1614%
18	30%	2134%
19	45%	3139%
20	-8%	2881%

Recall that the TRR is an unprofitable endeavor for the insurance company when the capital markets dramatically outperform expectations. Table 6 shows the single worst scenario (out of the 10,000 stochastic scenarios) for the present value of TRR claims. The first column shows the projection year, 1 through 20. The second column displays the annual net return to the policyholder's account value. The third column is the cumulative account value return since policy inception. A quick look at the table shows that the 10^{th} year cumulative return is almost 700% and the cumulative return by the 20^{th} year is nearly 3,000%. If a policyholder put $100 into this VA product at time zero, and he doesn't lapse or die, in the 10th year he would have $800 in his account. If that person then dies in year 10, the insurance company has to pay the beneficiary an add-on benefit equal to $280 (40% times the gain of $700). The beneficiary thus receives $1,080 in death benefit proceeds, but the insurer only has $800 in account value to release. The insurer's general account is thus on the hook for that $280 net amount at risk. (Note: in practice most TRR riders have a cap on the maximum gross-up amount.)

If at the start of this scenario an insurer sold a $1 billion block of new VA business with the TRR that I described above, the present value of the net amount at risk paid by the insurer to cover TRR claims would be $84 million.

The key thing to note in this scenario is that there are absolutely no GMAB claims paid by an insurer who issued GMAB products instead of TRR products. At no point during the projection period is the cumulative return less than 10 percent (which was the guaranteed growth promised by the insurer). It's never even close to 10% The GMAB writer made a lot of money in this scenario, while the TRR writer lost a lot of money.

Table 7

Year	Annual Return	Cumulative Return
1	8%	8%
2	-13%	-6%
3	-20%	-25%
4	-17%	-38%
5	-6%	-41%
6	-19%	-52%
7	7%	-49%
8	-9%	-54%
9	21%	-44%
10	-2%	-45%
11	-2%	-46%
12	1%	-46%
13	5%	-43%
14	-3%	-45%
15	-7%	-48%
16	11%	-43%
17	7%	-39%
18	11%	-32%
19	7%	-27%
20	16%	-15%

Now let's look at the worst case for the GMAB. Recall that the GMAB is an unprofitable endeavor for the insurance company when the markets dramatically underperform expectations. Table 7 shows the single worst scenario for the present value of GMAB claims. In this scenario, the markets dramatically under-perform expectations. The key projection year here is the 7th year (recall our design was a 110% guarantee on the account value on the 7th policy anniversary). The cumulative return in on the 7th anniversary is -49%. If a policyholder had originally deposited $100, and persisted to year seven, he now has $51 in the account value. The insurer has to step up to the plate and make them whole to $110 guarantee. The insurer must kick in $59 to get the policyholder's account value up to $110. To compound the insurer's troubles, this is also a scenario in which your revenue stream is down just because the account values are down, so it is really an unprofitable scenario. Selling $1 billion dollars of GMAB new business at the start of this scenario results in a present value of net amount at risk paid equal to $166 million.

How would a TRR product have performed over the scenario in Table 7? You'll notice that there is only one projection year where the cumulative growth is greater than zero, the first year. For any policyholders who die in the first year, the insurer would be on the hook for a fairly small TRR claim. After the first projection year the TRR is never again in-the–money (i.e. there is no taxable gain since the original premium exceeds the account value). The TRR writer made a lot of money in this scenario, while the GMAB writer lost a lot of money. The moral of Tables 6 and 7 is to illustrate that the scenarios that sink the TRR are scenarios in which the GMAB performs quite well and visa versa.

This leads us to explore a potential opportunity to exploit a natural hedge. The natural hedge is to use the liabilities themselves (the TRR and GMAB) as our hedging instruments. They have opposite risk profiles and opposite delta positions. The insurer is taking a negative delta position with the TRR and is taking a positive delta position on the GMAB. Now that we know that the delta positions are in opposite directions, we are ready to do something with that information. First we need to quantify the magnitude of the positive and negative delta positions. In this case study, the positive delta of the GMAB was about four times as large as the negative delta of the TRR. This indicates that the value of GMAB to the insurer is a lot more sensitive to changes in market movements than the TRR.

This 4-to-1 ratio tells us that if we want the present value of guaranteed benefit claims to be immunized against changes in account value, we will need to get an 80/20 mix between TRR and GMAB sales. If we can achieve an 80/20 mix, our delta position on the entire block will be zero, so we will be immunized. We won't care from a claims risk perspective if the account value goes up or down. It is important to note that saying that you are indifferent to market movements is sort of misleading because to maximize your revenue stream (primarily from M&E fees), which are the main profit drivers for VAs, you would want the markets to go up. However from a pure claims risk perspective, you don't care where the markets go if you immunize yourself.

Table 8
Cost in Basis Points

PERCENTILE

	TRR	GMAB	80/20
90	**35**	**95**	**31**
91	36	103	31
92	37	118	32
93	37	128	33
94	37	149	35
95	38	168	38
96	38	183	42
97	39	210	**47**
98	39	232	52
99	40	294	65
100	42	376	84
Mean	18	26	20
Vol	11	61	11

Let's assume that you are the pricing actuary and you want to issue these two products. Your corporate actuary tells you that you can do it, but you must price at the 90th percentile. You create Table 8 from the 10,000 stochastic scenarios you have run. Table 8 shows the break-even cost of the guarantee in basis points for various percentiles within the distribution. The break-even cost is the amount you would need to charge for the benefit such that the charges collected for the benefit exactly equaled claims paid on the benefit on a present value basis (calculated as the present value of net amount at risk claims divided by present value of account value). Taking your corporate actuary's desire to price at the 90th percentile into account, Table 8 says that if you are going to sell the TRR, you must charge 35 basis points, and if you are going to sell the GMAB, you must charge 95 basis points. Both of these fees are a little high compared to what is currently being charged for these benefits.

The power of natural hedging is shown in the third column of Table 8, labeled "80/20". The results in this column assume that an 80/20 mix is achieved (i.e. for every $100 of new business we sell, $80 of it is TRR and $20 of it is GMAB). The 90th percentile is now 31 basis points. The 31bps is obviously less than a simple multiplication of 80 percent times the 35bps for the stand-alone TRR and 20 percent times the 95bps for the stand-alone GMAB. That comes out to 47 basis points. If I was just going to use my stand-alone pricing and didn't run an aggregated run with an 80/20 split, I'd charge 35bps for the TRR (which is 80% of my sales) and 95bps for the GMAB (which is 20%) and on average I'd collect 47bps.

In the far right-hand column, 47 basis points is at the 97th percentile. If we charged this 47bps, we would actually be charging much more than our corporate actuary has told us we have to charge, which is fine if you can get it away with it from a competitive standpoint. Alternatively, you could exploit the power of risk diversification to become more competitive. You essentially have 16bps to play with, the 47bps less the required 90th percentile charge of 31bps. You can lower both charges by 16 basis points, knock off 16 basis points from the 35 and knock off 16 from the 95. Alternatively, you can get creative and allocate the 16 basis points however you want between the TRR and the GMAB. Let's say I was new to the GMAB market and I wanted to buy some market share there. I could allocate all of the 16 basis points to the GMAB line (which is 20 percent of my portfolio) and knock the 95bps charge down to 15bps (16bps divided by 20% is the 80bps reduction).

Alternatively, if I think I can sell these products at the stand-alone pricing rates, I could just go out to market with these prices (35bps and 95bps). Then I would be protected at the 97th percentile, so I'll be taking less risk than had I priced at the 90th percentile.

Another very interesting thing to note is how the downside exposure has gone down dramatically. From the GMAB results in Table 8, look at 100th percentile value of 376bps. I can never sell this product at 376bps, so I'd always be exposed to painful loses in the 100th percentile. However, if you look at the column for the 80/20 split you will see that we have dramatically reduced our downside exposure. In the 100th scenario we would have needed to charge 84bps to break-even. If we price at the 90th percentile our biggest shortfall is the difference between the 31bps and the 84bps. In our worst scenario, we would only be short 53bps from a claims perspective. This 53bps shortfall could probably be significantly covered by the profits from the rest of the VA contract.

In conclusion, we have used power of offsetting risks to reduce the risk profile of our VA product line.

From a practical perspective, this is not an easy thing to do. I've laid out four challenges that I think you may encounter if you were going to try and implement something like this.

The first is balancing the mix of sales. The 80/20 sounds nice, but how do you do it? You can have sale specials. If one's not selling fast enough, you can offer bonuses or DCA specials or enhance the compensation to the agent. If one's selling too well, you may have to do something too slow it down (you might even have to pull the product from the market). You don't have to do this naturally, by managing sales. You can do it artificially through the use of derivatives or reinsurance. You could do something similar to what I'm talking about and then also bring in what Charles was showing you, using equity derivatives, to manage your delta position. Let's say my delta position was way off because I had too much GMAB on my books. I could enter into the equity markets and sell delta (short futures) and bring my portfolio back to a delta neutral state. Alternatively, I could reinsure some of my GMAB exposure. If I was selling too much GMAB, and I didn't want to pull out of the market, I could start seeding some of that away.

The second challenge deals with the fact that delta (like the rest of the Greeks) can change over time as economic conditions change. Therefore, you really have to be able to model your in-force block and your anticipated new sales accurately, frequently, and quickly. You don't want to waste a week doing a run and then realize that you're delta position is way off.

Probably the hardest challenge is the third one, and that is developing marketable products that offer natural hedges. We as actuaries have got to get creative and think of ways we can do this. The TRR and the put option guaranteed benefits (like a GMAB) are the obvious natural hedge. Another obvious one is put option-type guaranteed benefits in these VA contracts, such as rollups, return of premium, and then the call option type, equity-indexed annuities (EIA) that people sold a bunch of back in the early '90s. To hedge these EIA products, insurers are buying a lot of delta, and to hedge some of the put option guarantees, you'll be selling delta. Why not get the two lines together and net it out, thus entering into fewer equity options and save the transaction costs?

The fourth challenge is that you need a sophisticated modeling platform to be able to stochastically project your in force and your new blocks of various lines of business. If you start bringing other lines of business (such as EIAs), then you will need more than just a VA model. The more creative you get, the more your modeling needs expand.

Wrapping up, I'd like to share with you a few of the broad observations that I have made regarding the issue of VA guaranteed benefits. I think it is essential that the product development process is not done in a vacuum. For example, let's say you're looking at just pricing the TRR and your counterpart who works for the same company has been given the challenge of pricing the GMAB. If you two don't talk to each other, neither of you will be getting the full picture and you are not going to be able to optimally exploit the benefit of the diversification. When you analyze the risk profile of these two benefits separately, both may be unacceptable for your risk tolerance and you may not end up issuing either product. However, when you analyze the two products together, you may be able to launch these products at acceptable risk and profit levels. The benefits of risk diversification can allow you to launch products that you typically may not have launched.

Another key thing I think is that you can't price under a follow-the-leader mindset. I think when a lot of these benefits first came out, everyone was trying to catch up to the first guy who went to market. All of the designs and prices seemed to be about the same. If your competitors are doing analysis like this and using the power of risk diversification to allow them to offer a low price on benefit and you are not, you are taking on more risk than you are being compensated for if you match their price.

Offering both of these products may allow you to offer more aggressive prices. Exploiting natural hedges will lower your derivatives-based hedge costs and/or your reinsurance costs because you will not need to hedge/reinsure as much business. Another key is if your reinsurer assumes both of these types of products and you are sending them both types, you want to make sure that you've done this type of analysis. That way, you know what you should be paying for your reinsurance because there is the netting effect there.

One final point I'd like to make is that I am a big fan of using a three-legged approach to VA guaranteed benefits risk management. Using the three risk mitigation techniques we're talking about here—natural hedging, derivatives and reinsurance—in tandem, I think, is the optimal strategy for an insurer to protect itself from VA guaranteed benefit claims risk. You don't want to just live or die by one strategy.

MS. HARRINGTON: I'm going to be talking about the reinsurance market from my perspective. As I mentioned, I've been involved in the reinsurance of variable annuity guaranteed benefits for the last seven years. I'm going to talk about the history of the market, the current state of the market, and what to expect if you're going to look for reinsurance for your company. Hopefully, after this session, you'll have some ideas on where to start and what to think about when you're looking for reinsurance.

Reinsurance of guaranteed benefits in variable annuities began with the guaranteed minimum death benefit in the early 1990s. It was a very competitive market at that time. It was possible to find reinsurance on original terms, meaning that you could find first-dollar reinsurance with no cap on the benefit at premium rates that were fairly close to your retail prices. Swiss Re and Cigna Re were both very active in this marketplace. There were other companies involved as well. As we moved through the 1990s, the guaranteed benefits became increasingly popular on the retail side. They became richer in design, and living benefits were added as well. Things were going along just fine from the ceding company's perspective. Then in about 1998, both Swiss Re and Cigna Re pulled out of the market. They exited for different reasons.

Having been at Cigna Re at the time, I know that we had no idea what was going on at Swiss Re, so it was as much of a surprise to us as it was to the marketplace. It was unfortunate that the timing was such that both companies left at the same time, because I think it gave people the impression that something was wrong with the product or with how we were reinsuring it. That really wasn't the case.

One company exited the market because a new president came on board who had never liked this risk and did not want his company to write it. For the other company, a combination of things relating to how much business they had already written as well as senior management's view of the worst-case scenario of what could happen to their business led to its exit. When they combined the volume of business they had written with their assessment of the worst-case scenario, they realized they were getting too close to the maximum they would be comfortable writing.

As I'm sure you've heard in many of the sessions, this is a seller's market right now. It is very tight. Depending on which sessions you went to, you may have come out with the impression that you might as well not even bother to look for reinsurance—it will be impossible to find. You may have left other sessions feeling somewhat better, hearing that there is a market for reinsurance out there. And I can attest to the fact that there is. Over the last six months, we've placed several treaties. Even in the last month we've placed some treaties with reinsurers covering the death benefit, the income benefit, the accumulation benefit, and the tax relief benefit. So, it is possible to get the reinsurance.

AXA Re is the main player domestically, but there are also offshore companies interested in reinsuring these risks. Some of these offshore companies are property and casualty (P&C) carriers, and they view this risk differently from life companies. They don't view it as something that they need to hedge in the capital markets. They view it more as a catastrophic risk, which they're very familiar with. A lot of the business they write covers catastrophic risks related to natural events such as earthquakes, hurricanes, and other disasters. By reinsuring stock market performance, they're adding another event that is, for the most part, uncorrelated to the risks that they already have in their portfolio. Because of that, they're willing to take some of this risk. This doesn't mean that they're aggressively seeking this business, but they do see the value of doing it in certain circumstances.

Today's reinsurers are willing to take risks that make sense, and I put the emphasis on "that make sense." It makes sense to them if there's significant volume potential. They don't want to spend a lot of time on something that's not going to generate enough income for them.

If your business is not that big and is not expected to grow quickly, you still have options for obtaining reinsurance. The opportunity to cede other business with the reinsurer can help you gain reinsurance for your guaranteed benefits. In fact, this is where I would start if I were charged with finding reinsurance for the death benefits or income benefits under my variable annuity products. I would look at who my current reinsurers are. You have a lot of leverage.

For example, look at who's reinsuring your life benefits today, whether group or individual life, or who's reinsuring your fixed annuity products today. You can ask these companies for help reinsuring your variable annuity products. If they don't want to or can't help you and you're willing to move your reinsurance, you can use that as leverage. You can bring the life or fixed annuity business to another reinsurer who is very happy to have this business and is also willing to accept the risk from the variable annuity block.

If you're unable to find reinsurance at all, it may be an indication that your product isn't viable. Maybe there's something in your design that just doesn't work or there's the potential for too much antiselection. If you make a good effort to find reinsurance and are totally unsuccessful, I would definitely go back and take a look at the product.

What reinsurers are looking for today in their treaties, whether it's a first-dollar treaty or a stop-loss treaty, are limits. They want the worst case specified in the treaty. By this, I mean a maximum claim limit. It is often expressed as a dollar amount per $1 billion of retail annuity premium. This is a very important number to have in there. It drives the premium. It drives the reinsurer's capacity. They look at how much business they've written and determine their claim exposure under the worst case. If this sum gets too high, it may limit their ability to write more of this business.

Reinsurers also want a finite term; 15-year and 20-year terms are quite common. If you felt that your particular block of business is going to persist, that you're going to have a significant amount of that business in force 25 or 30 years from now, certainly you can talk to the reinsurance companies about extending it. Both the claim limit and the finite term are negotiable items. Of course, everything has a price associated with it. The higher the claim limit, the higher the reinsurance price, even if the likelihood of reaching the limit is almost non-existent.

The reinsurers are looking for premium commensurate with the risks. They have very sophisticated stochastic modeling tools that enable them to analyze and quantify these risks. They don't want any surprises. They're certainly willing to pay the claims when account values fall, but they don't want to discover that their mortality assumptions or lapse assumptions were way off base. They're very concerned with getting accurate demographic and exposure information from the ceding company to help set the assumptions for their modeling. The more comfortable the reinsurer is with the assumptions, the more likely they are to accept the risk.

There's a lot of talk about how you can't reinsure the tail. That's only partially true. The tail of the risk does drive the premium. That's what the reinsurers use to set their premium levels. The deductible can be set at any level, which allows the ceding company to retain as much of the up-front risk as they wish. The excess over the deductible, the tail, is then reinsured subject to a claim maximum. Companies would certainly love to reinsure the entire tail with no limit on claims, since most of the product designs don't have a cap on the retail product. I've had many companies comment that they think its the job of the reinsurer to take the tail of the risk. In today's environment particularly, reinsurers are not willing to take this risk on an unlimited basis. However, they are flexible with where the limit is set.

The ceding company has to establish the worst-case investment scenario for which it wants protection. While the actuary will most likely recommend the scenario, senior management should buy into it. You don't want to be so conservative in setting this that you pay for a lot more reinsurance than you need, and you don't want to set your worst case so low that you don't get the claim coverage that you need. You need to set a meaningful scenario. We have found that tying the worst-case scenario to the Great Depression works well for many companies. Typically, a 50 percent drop in the equity return, followed by zero percent recovery for the next 15 to 20 years is a fair proxy for the Great Depression. Once you have agreement on the worst-case scenario for your company, you can determine the claim limit needed to ensure that claims under this scenario would be fully covered by the reinsurer. This is done by modeling the block of business to be reinsured using the worst-case investment returns plus realistic demographic, mortality, and lapse assumptions.

Here are some numbers to give you a feel for what these claim limits look like by product:

GMDB - $101 million ($57 PV)
GMIB - $217 million ($99 PV)
GMAB - $306 million ($176 PV)
GPAF - $273 million ($144 PV)
These numbers relate to $1 billion of retail annuity premium. They were derived for one particular client. At these levels, the company was confident it would be protected even if we experienced another Great Depression. You'll see the GMDB is at the lowest with $101 million on an undiscounted basis, or $57 million, discounted. The GMAB is the highest, at $306 million per billion. For this client, without a doubt, its GMAB would be the most expensive reinsurance to obtain.

One of the things that has changed over the years in seeking reinsurance is how you go about doing it. You really need to do your homework in advance to know the claim limit you want and how much premium you're willing to pay. You have to be prepared to share this with potential reinsurers. When I was at Cigna Re, companies would submit open-ended proposals in which we would determine the terms and premium. That approach doesn't work today because (1) the market is tight, and (2) reinsurers have limited time and resources to spend. The analysis of these risks takes a long time, and they want to focus on opportunities that are going to be serious, that have a high likelihood of generating reinsurance. When the ceding company shares its expectations upfront, the reinsurer can quickly tell if they are realistic and if it makes sense to invest time analyzing the proposal.

I think some reinsurers got burned in the past, and that's how we evolved to this point. They would invest considerable time analyzing the risk, only to find that the reinsurance premium was well above the ceding company's expectations and that they did not want to buy reinsurance at this price. The reinsurer lost valuable time it could have spent on other business opportunities.

Doing the analysis up front and sharing expectations regarding the terms and price has worked well for our clients over the last few years. Not only do the reinsurers like this approach, ceding companies find it expedites the reinsurance process internally.

In doing the analysis, it is important to understand how the reinsurer views the risk. This often differs from the ceding company's view. Assumptions are set differently since the reinsurer only has the guaranteed benefits risk, where the ceding company has the entire variable annuity. Consider lapse assumptions. Lower lapse assumptions are conservative when looking at just the guaranteed benefits since more contracts will be in force to experience a claim. Higher lapse assumptions are conservative when viewing the full variable annuity since less contracts will be in force to pay the contract fees.

Reinsurers tend to use a standard set of mortality assumptions in their pricing models. However, they will alter these if a ceding company has performed an internal mortality study that it is willing to share. This can definitely help you. We've had several clients whose mortality experience was significantly better than the standard mortality assumption that many of the reinsurers were using. By providing the study to the reinsurer, they were able to obtain better terms.

What else should you think about? I can just reiterate what Dave had mentioned, that you really want to get the reinsurers involved early in the process, in the product design process, because then there are no surprises. The reinsurer can help you with the retail design and the pricing, and you can avoid delays later on if you get them involved early.

Finally, something that's been very effective for us is leveraging the natural hedging property of the tax relief benefit. Natural hedging works for the reinsurers; it improves their overall risk portfolio, just like it does for the ceding company. But it can often be used more effectively by the reinsurer since they have better control over their product mix. They can assume more reinsurance of one product versus another to try to get the optimum balance. Product mix is more difficult to manage on the retail side.

Reinsuring the tax relief benefit can help in two ways. It can help you obtain reinsurance capacity for the other benefits that may be more difficult to place, particularly the accumulation benefit or the income benefit. It also may enable you to improve your overall reinsurance terms.

MR. RUSSELL A. OSBORN: I just have a quick question for Charles. In your Greek hedging example, you hedged vega and gamma first and then delta hedged. Was that just for illustrative purposes, or would you normally do that in practice? And, if so, why?

MR. GILBERT: I think there are a number of different hedging strategies that you can use. If you are going to hedge delta/gamma/vega, typically you would hedge the gamma and vega exposure first, using the two equations, two unknowns. You do that because then you're going to be left with the delta hedge. The delta exposure is easily hedged using futures, because futures have a delta of one and a gamma and vega exposure of zero. So it’s very easy, once you’ve hedged your gamma and vega exposure, to just get rid of the residual delta exposure.

Chart 1

Guarantee Cost Example - Risk Profile of GMAB Hedged and Unhedged

3

Chart 2

Risk Management Options

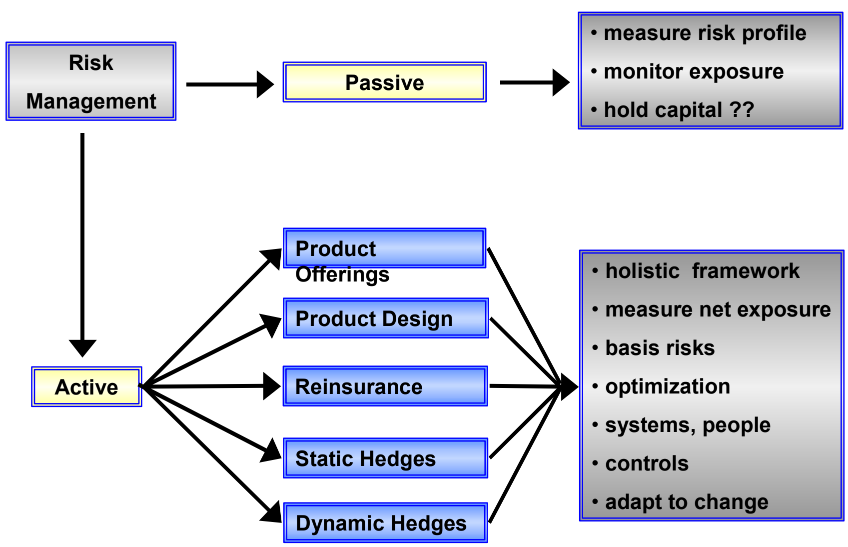

7

Chart 3

The Dynamic Hedging Process

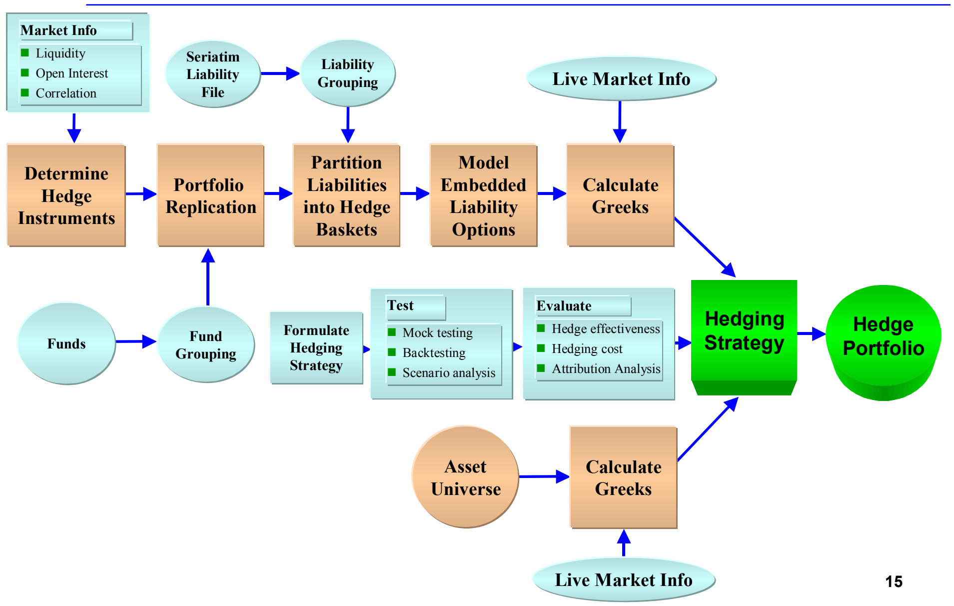

15

Chart 4

Key Practical Issues

- Volatility
  - static hedges remove exposure to future volatility
  - delta hedging leaves you exposed to actual volatility
  - static hedges are priced in the capital markets
  - market makers charge an implicit volatility premium (6%+ vol)
  - delta hedging attempts to remove the market premium

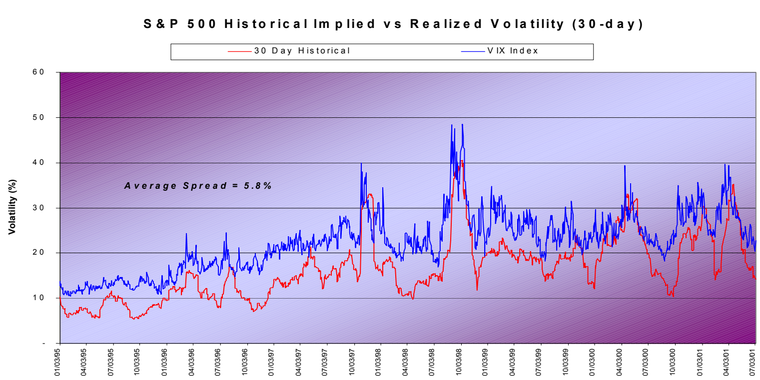

25

Chart 5

Key Practical Issues

- Trading frequency vs. transaction costs
  - more frequent re-balancing improves hedging performance
  - more frequent re-balancing increases transaction costs
  - long term nature of liability guarantees
  - strategy must deal with normal and turbulent markets