RECORD OF SOCIETY OF ACTUARIES
1995 VOL. 21 NO. 1

VARIABLE ANNUITY GUARANTEED MINIMUM DEATH BENEFITS

Moderator: GARTH A. BERNARD, SR.
Panelists: STEVEN P. HABEGGER
G. THOMAS MITCHELL
JOHN SLATER, JR.*
Recorder: GARTH A. BERNARD, SR.

Panelists will discuss the structure, cost, and reserving practices of variable annuity guaranteed minimum death benefits.

MR. GARTH A. BERNARD, SR.: I'm head of retail product development at Providian Capital Management. And I'm not going to spend a lot of time talking about my interest in guaranteed minimum death benefits on variable annuities because we have an excellent panel here.

I'm going to introduce our panelists in the reverse order in which they're speaking. John Slater is vice president at UBS Securities. He spent six months in the industry on the sales side with Northwestern Mutual Life, but for the past ten years, he has been structuring risk management solutions for stock and bank portfolios. I've asked John to speak from the perspective of an investment banker by talking about guaranteed death benefits; they actually tend to look somewhat like put options on securities.

Our second speaker will be Steve Habegger who is director of Individual Life Insurance Products for Lincoln National Reinsurance Companies. He's responsible for the pricing and product work for East and West Coast reinsurance clients. Prior to working in reinsurance, Steve did annuity product work for Lincoln's direct site. So I've asked Steve to talk about this topic from the perspective of a reinsurer that would be interested in assuming these types of risks.

Our first speaker is Tom Mitchell. Tom is an independent consulting actuary with his own firm, Aurora Consulting, located in St. Louis. Tom has been a Fellow of the Society of Actuaries since 1966 and is also a member of the American Academy of Actuaries. Prior to starting his own actuarial consulting practice in January 1991, he was president of Charter National Life Insurance Company. So without further ado, I'll ask Tom, who has had substantial experience from the variable side when he developed seven products at four different companies to speak. His perspective will be from the direct writer's side.

MR. G. THOMAS MITCHELL: By way of introducing this subject, I once had the embarrassment of representing a company that put out a brochure, not a variable annuity, but on a life insurance policy to be sold with individual retirement accounts (IRAs), that had a sales pitch that went as follows: "We can help you meet your retirement objectives, even if you should happen to die before retirement."

---

*Mr. Slater, Jr., not a member of the sponsoring organizations, is Vice President of UBS Securities Inc. in New York, NY.

I want to start out by going through the working parts of the variable annuity minimum death benefit design, because the vocabularies used are different in each situation. So you'll at least know the words I use to describe it. The classic formula is guaranteed minimum return of premium on death, prior to annuitization and net of earlier partial withdrawals. Nowadays, we find features where you can add what I call a reset to it. This means that every five to seven policy anniversaries, the minimum is increased (but it would never be decreased) if the account balance exceeds the minimum death benefit on that anniversary date.

Another twist is to credit the minimum death benefit at a stated rate of interest, either at the simple or compound interest rate. Being inventive, we can then combine the interest growth rate or the reset features to get an even fancier benefit. Either feature can be capped or limited according to issue ages, attained ages, percentage increases, or anything else you'd like to devise. The benefit is customarily charged out of the mortality and expense risk charge, as basis points on the account balance. But it can also be financed by a premium load which could vary by issue age.

I'd like to look now at what's prevalent in the marketplace. We surveyed the top 25 products of 1994, listed by VARDS, and determined what provisions are currently in place for 22 out of last year's 25 products (see Table 1). One company actually doesn't even seem to have a minimum death benefit, but that's the exception. The most common is the return of premium or the classic formula. We see that nearly a quarter of the marketplace has a reset of every five to seven years.

One of the latest things is to reset it every year, instead of just doing it every five or seven years. There are two products from the same company that do that. Products that increase at 5% are a little less popular. Finally, there are two products that have a combination feature. So there is some variety there. One of the things to notice is, although the trend is towards enhanced benefits, much business is still being sold on the old conventional formula.

TABLE 1
MINIMUM DEATH BENEFIT PROVISIONS
22 OF TOP 25 VARIABLE ANNUITIES

Benefit Type	Number	Percentage of Sales
No Minimum	1	5%
Return of Premium	10	39
5–7 Year Resets	5	24
Annual Resets	2	14
At 5% Interest	2	12
Better of: 4–5% Interest or 7 Year Reset	2	6
Total	22	100

It's getting rather difficult to top somebody else's benefits and have the best in the marketplace, because we've seen a reset every anniversary. I've also seen a product where the compound interest rate growth was 7%, and current regulatory constraints are going to make it unlikely that these can be topped.

What are the marketing dynamics behind offering the enhanced definitions? I think the first dynamic is the defense against replacement. With the return of premium design, most individual contracts, most of the time, will be out of the money on the minimum death benefit. After the surrender charge wears off, an agent can encourage a client to increase the death benefit by rolling the policy. That gives somebody who wants to twist the policy something to talk about. The best design or best defense against this is the periodic reset, because by using that design, you catch up on the death benefit every year, five years, or seven years and it never gets too far out of line.

I think the second dynamic is simple—features competition. The enhancements are bells and whistles for the wholesalers or the marketing vice president or whomever is out there pushing the product. There's a lot of competition for shelf space. I think an agent is probably slightly interested, not wanting to lose a case to somebody who has a fancier benefit. But I seriously question whether the trend to stronger death benefits is really consumer driven.

Let me focus now on what these benefits are likely to cost. First, there's my methodology. I ran 1,000 scenarios on the U.S. equity stock market, using the monthly Standard & Poor's 500 Index total returns, including dividends for 1950–93, as a basis for establishing stochastic distribution. What I used is what is called a symmetric stable paretian distribution, which does a better job than the normal distribution in modeling the frequency of prices and large market runups. An article in last October's *Risk & Rewards* presents more details. What I've done is to present value death benefit costs for a no-load, single-premium annuity in terms of dollars and cents on $100 of single premium.

Let's take a look at some of the results. Chart 1 shows numbers that are averages of over 1,000 scenarios. Based on historic performance, the simple formula, on average, doesn't cost a great deal. Actually, $.08 was the number that I came up with. Not very surprisingly, it's highly correlated with age. It's obvious, but interesting to note, that the charges levied for the death benefit don't usually vary by age.

Now, we reset the benefit on different anniversaries (Chart 2). We get a far different picture. The first example is labeled infinity. That means you never reset it. That's the $.08 that we saw before. But if you reset at any frequency, the average cost of the benefit rises significantly. The benefit just doesn't wear out with age. It's still in the money at the higher ages. It goes up for five or three years, and for one year it resets; it's much higher at $1.77. One of the reasons is, every year that the market goes down instead of up, there's a cost. The thing literally never wears out, so there is a very significant cost with the resets.

How does the cost vary by growth with interest rates (Chart 3)? With 7% interest, I have $0.34. So, it's a large multiple of the straight benefit, but it's not nearly as costly as the reset methodology.

CHART 1
COST BY ISSUE AGE

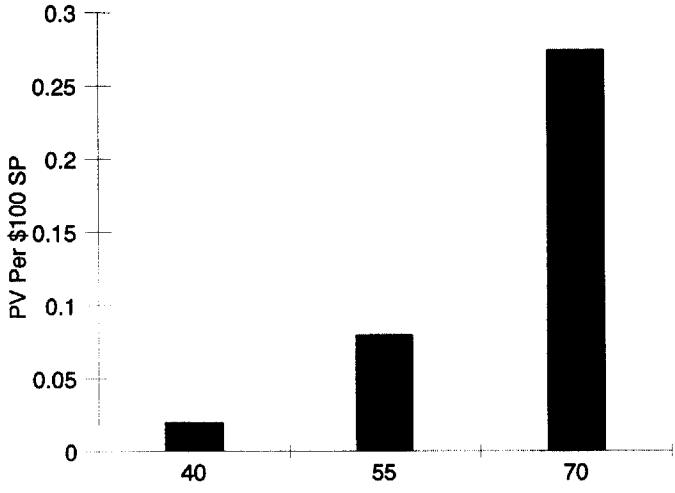

CHART 2
COST BY RESET FREQUENCY IN YEARS

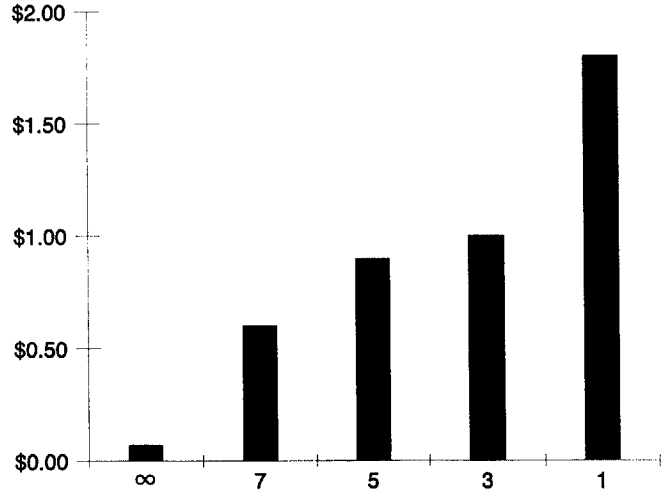

CHART 3
COST BY INTEREST RATE

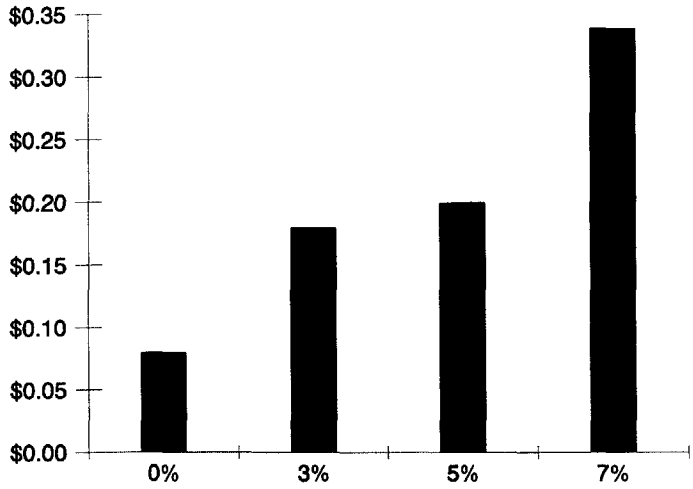

Now, that's not all of the story. There's much more to it. We need to look at the distribution of the costs, and Chart 4 shows a rather odd distribution. If you run 1,000 scenarios, you get, on a simple benefit, no cost whatsoever for a little over 80% of the time. This is on an annual model and there's some simplifications in practice obviously. You always have at least some costs, but the formal distribution is zero about 80% of the time. On the other hand, if you go to the worst percentile and average the worst percentile, you get $2.40, which is a pretty good bite out of variable annuity profit margins. Within that last 1%, obviously, there's some that are better and some that are worse. These projections make no provisions for the problems that you might have with antiselection in a bad scenario, in terms of mortality. You can have antiselection in terms of lapses. You can have antiselection in terms of partial withdrawals. Real results would undoubtedly be very much worse. So we have something that most of the time doesn't cost very much and sometimes costs a great deal of money. There's a great incentive for reinsurance or hedging and that's what we're going to hear about from the other two speakers.

What are some of the things you should think about in the design process? First is protection from replacement. You can probably meet that best with a reset design. Distinctiveness in the marketplace is a little bit of a problem. There may be some new twists on formulas that you can come up with, but competing on the amount of benefit is going to be difficult. A third factor you should keep in mind is perceived value and understandability to the customer. I think many of the advance designs are rather difficult for the public to understand.

CHART 4
DISTRIBUTION OF COSTS

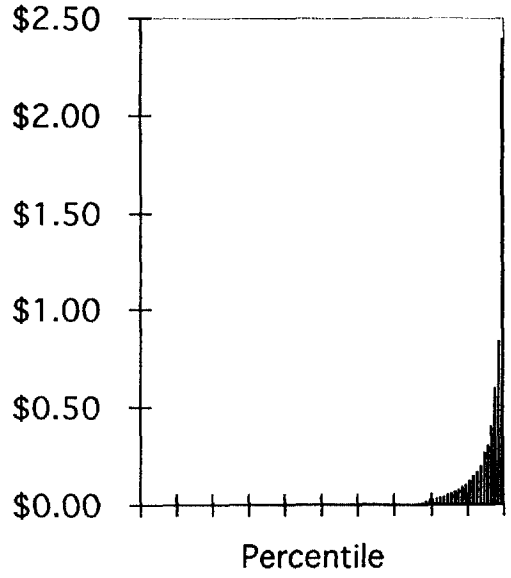

There are two elements of cost control. One is to control the average cost of the benefit, and that, obviously, can be done with caps on the design or restraint in the design. The other problem is what do you do about the fact that this is very costly in a small percentage of possible futures? Besides restraint and wisdom in the product design, it's possible that marketing strategy might help. The other two tools are to use reinsurance or to go directly through hedging.

I'm going to leave you with a little more sense of urgency about some of the catastrophic scenarios. If we get into a down market, the customers can often make very poor timing decisions on transfers, lapses can go up, and you can have a deterioration in the mortality experience. Another thing that can happen is many companies are wide open on the partial withdrawal provisions.

I want to tell you a true story. This is coming out of a direct response variable annuity solicitation. A respondent calls in and says, "I have a question for you. You mean to say that if I give you $100,000 and the market goes down 20%, I'll only have $80,000? Also, is your minimum account size $2,500? Do you mean to tell me that I can make a partial withdrawal? I can get $77,500 out?" "Yes (gulp)." "Have $2,500 left and then if I hold that until I die, you're going to give me at least $20,000?" And the answer was, "Yes. That's what we mean (second gulp)." There is some potential for antiselection. When the stock market is going well, this doesn't look like a very real risk. But if we were in a severe market downturn, for a year or two, and Jane Bryant Quinn wrote an article titled, "Your variable annuity policy may be a dog, but check the death benefit before you surrender," and goes on to explain further in the article about making partial withdrawals; it could have an effect.

Let me turn now to an example of how much the benefit can cost (Chart 5). We made up an example. This is with a formula that includes a 7% increase each year. One of the things to notice here in this example is that the market plummets quite a ways, and then it recovers pretty well, but because of the 7% increase, the death benefit is in the money for a substantial period of time. Then it actually gets out of the money and goes back in again. In this scenario, which is on a 70-year-old male, the annual cost to the company per $100 of original premium (Chart 6), varied during the first seven years between $.50 and for the worst year it is almost $1.50. That, friends, is a pretty good bite out of anybody's variable annuity profit margin. I don't think there would be anything left there.

Actually this is not a contrived example. This is what would have happened if you'd written the policy with this death benefit formula on a 70-year-old male, as of January 1, 1973. It shows what actually happened in the marketplace and what the resulting costs were. There are, in conclusion, some very real risks, which is why we have two more speakers to talk about what you can do about it.

CHART 5
DOWN SCENARIO

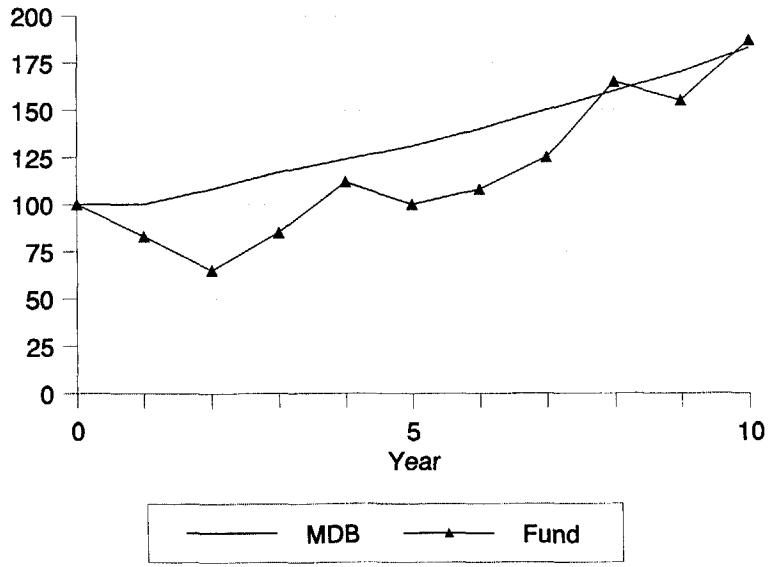

CHART 6
DOWN SCENARIO
ANNUAL COST PER $100 SP

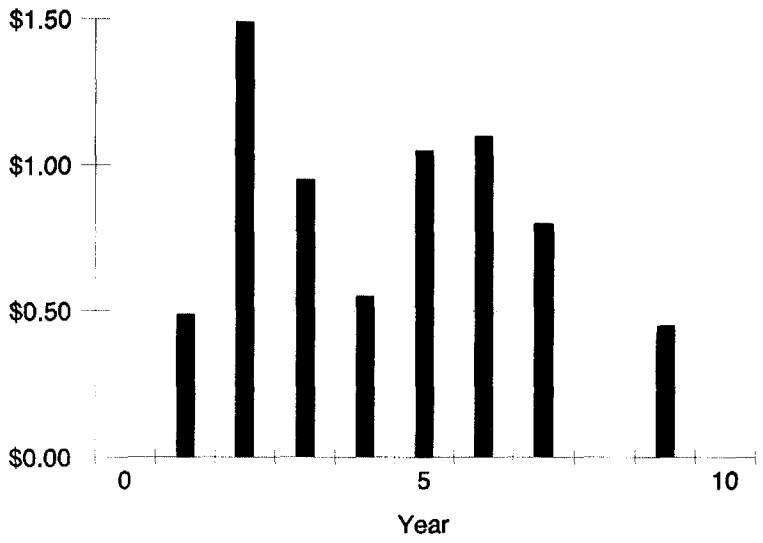

MR. STEVEN P. HABEGGER: The reinsurance industry is in the business of taking on insurance risks. The risk transfers we typically participate in have three broad characteristics. They are diversifiable: the more of a certain type of risk we take on, the more the variability in results decreases. Second is limiting antiselection: we try to minimize the possibility for antiselection. Third, there's partnership: in most cases we want the ceding company to be our partner in the risk-sharing relationship.

RISK IDENTIFICATION

From our perspective, the starting point is to accurately identify and measure the risks that are present in any deal we are considering. There are three main risks associated with guaranteed minimum death benefits (GMDBs):

- market risk: the investments underperform the guarantee offered
- mortality risk: more policyholders than expected die
- product design risk: something in the contract allows policyholders and/or sellers to select against you.

I will discuss each of these risks, draw some parallels to life reinsurance where appropriate, and suggest a structure for handling each of them.

MARKET RISK

Market Fluctuation

The most apparent risk is simple market fluctuation. We all know that the stock market is volatile and may or may not outperform the guarantees offered on GMDBs. Despite the higher anticipated performance from equities, insurance companies limit their investment in stocks because of the potential risks to surplus levels, ratings, and solvency. Quite

simply, by writing GMDB coverage, insurance companies are placing a bet on the stock market. A typical insurance operation is not equipped to take on this risk. This risk is nondiversifiable without hedging. That is, when a number of ceding companies pass this risk on to a reinsurer, the absolute volatility the reinsurer can expect actually increases. This is a risk that we must get out to the financial markets where there are instruments, investors, and speculators who are experts at taking on and managing this risk.

Fund Manager Performance

There is a risk that a fund manager underperforms the general market. When we price the market risk or purchase a hedge for the risk, some assumption is made as to the expected return and volatility of that return. The "good" manager is one who closely mirrors whatever we expected of him, even when that means losses. I think a parallel can be drawn here to life underwriting. Reinsurers may need to develop expertise at underwriting the fund managers. How good are they? How do they perform compared to the overall market during upcycles and during downcycles? This is a risk that can probably be shared by the reinsurer and the ceding company.

Fund Distribution Risk

The market risk associated with GMDBs depends on how much money is in each fund type: small cap stock, large cap stocks, bonds, or money market funds. There is no good parallel to life insurance risk here because with life insurance, policyholders are charged for the extra risk they bring to the pool. Thus far, the annuity marketplace has not had a product that varies the price by fund allocation. The best place for the risk may be the policyholder. Those policyholders who wish to invest in riskier funds, should probably pay their fair share. Until the market changes, reinsurers will be willing to share this risk to the direct company. How much can this risk impact the price of the GMDB? Chart 7 illustrates the difference in cost between a small cap and a blue chip stock fund.

CHART 7
FUND PARAMETERS—5% INTEREST GUARANTEE
EFFECT OF MARKET ASSUMPTION

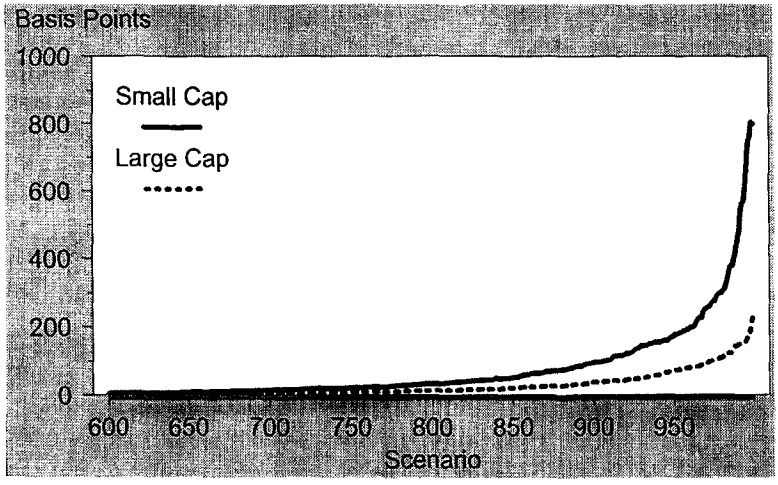

MORTALITY RISK

Mortality Fluctuation

The risk that more people die than expected has direct parallels to life reinsurance. This is a risk that life reinsurers are expert at handling. It would be ideal for the direct company to share in this risk to some extent. The alignment of interests is paramount in a risk-sharing arrangement.

There is a new twist on the mortality fluctuation with GMDBs in that the actual claim amount of the excess deaths depends on the market performance. Thus, you need to model a joint probability distribution in order to measure the mortality fluctuation risk.

Large Deposits

Large annuities can have great impacts on all three distribution risks I have discussed. When an extremely large deposit is made, not charging the annuitant accurately for the funds they have chosen, or for their age/sex contribution to mortality risk, can create serious problems in your pricing. Careful consideration needs to be given when dealing with large policies. A reinsurer who has accepted all the GMDB risk for a ceding company, may not even get to know that there is such a large policy. One possible solution is to limit the policy size that is allowed to have a GMDB. Another partial solution is for the ceding company to reinsure this single policy on a yearly renewable term (YRT) basis, purchasing reinsurance only when the guarantee is greater than the account value, and paying an age- and sex-distinct rate. This is not the optimal solution, because the ceding company has still retained all of the market risk. It would not be necessary to pay a fund-distinct rate in this case, because the reinsurance premium is due only when the net amount-at-risk is positive.

Distribution Risk

Payment of this benefit does require that the policyholder die. So, like life insurance, the cost depends on age and sex. In life reinsurance, we often take on age and sex distribution risk created by the ceding company's need to simplify administration. With GMDBs, the distribution risk is actually created at the policyholder level. To date the annuity marketplace has not differentiated the cost by age or sex. In addition, differentiation by sex may be limited in some markets. This is also a risk that may be shared between ceding companies and reinsurers.

Charts 8 and 9 illustrate the impact that varying the age distribution and sex distribution can have on the cost.

PRODUCT RISK

Withdrawals

The last risk I want to discuss arises because of the way the benefit descriptions are worded. Currently, if a policyholder withdraws all but some minimum amount, only the guarantee level is reduced. This results in the amount at risk remaining constant. Table 2 illustrates the impact this can have and a couple of possible solutions.

CHART 8
AGE DISTRIBUTION
5% INTEREST GUARANTEE
EFFECT OF AGE ASSUMPTION

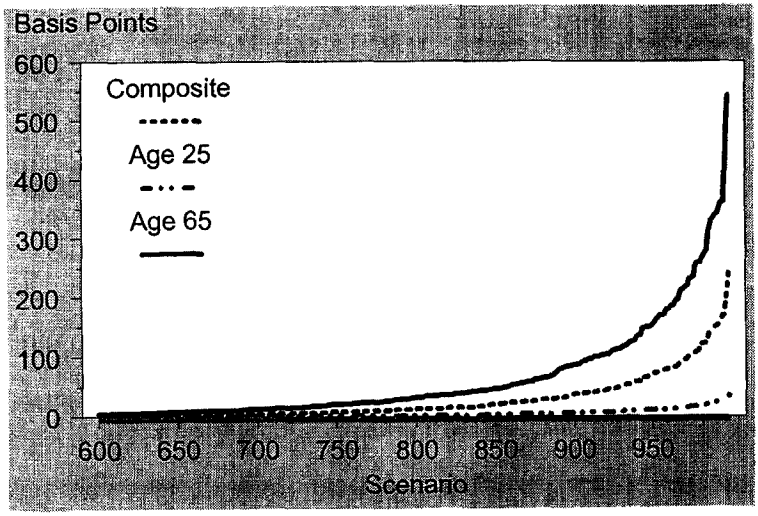

CHART 9
SEX DISTRIBUTION
5% INTEREST GUARANTEE
EFFECT OF SEX ASSUMPTION

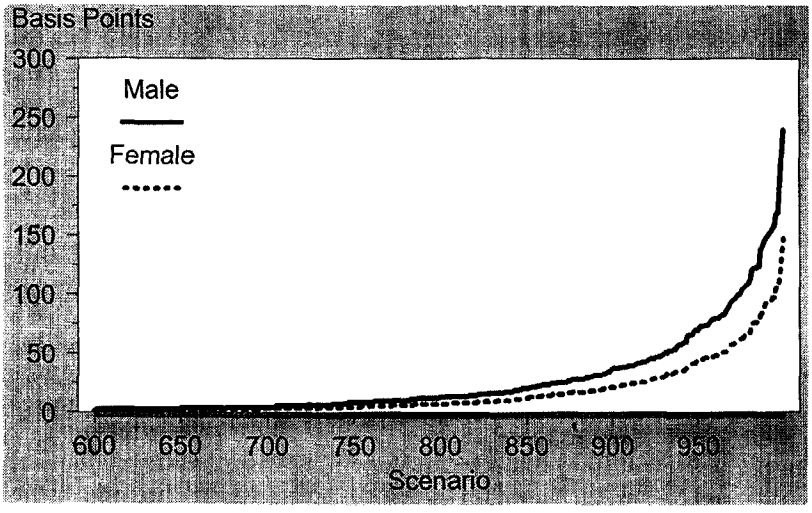

TABLE 2
WITHDRAWAL RISK EXAMPLE

	Absolute benefit	Absolute at risk	Proportional benefit
Before withdrawal			
Account value	$ 40,000	$ 40,000	$ 40,000
Guarantee	100,000	100,000	100,000
Amount at risk	60,000	60,000	60,000
Withdrawal	39,000	39,000	39,000
After withdrawal			
Account value	1,000	1,000	1,000
Guarantee	61,000	21,000	1,500
Amount at risk	60,000	20,000	500

The first column illustrates the current wording: the guarantee is reduced by the withdrawal. This results in the amount at risk remaining constant. The second column offers another option: the withdrawal is from the amount at risk to determine a new guarantee. This option still poses withdrawal risk to a ceding company. The final column illustrates a proportional reduction in the guarantee. The withdrawal was for over 97% of the current account value, so the guarantee is reduced by a proportional amount.

There are a number of problems with not reducing the guarantee proportionally:

1. You are giving away free life insurance: you are not getting any premium.
2. The account base is greatly depleted: you have no source of profits.
3. The probability of recovery is zero: 100% returns on stocks will not return the account value to its previous level.

This risk benefits the withdrawing policyholder to the detriment of those that stay in for the long haul and the insurer. This risk can be reduced by changing the benefit description or basing the premium on the guarantee value rather than the account value.

A PROPOSED STRUCTURE FOR SHARING THE RISK

Many ceding companies have desired to cede all of this risk to a reinsurer. While this is the simplest answer for the ceding company, unless the reinsurer is hedging the risk, the life insurance industry still has considerable market risk.

Limited Price Changes

We need to hedge the market risk out of our industry. When things go wrong here, people are going to be running for the hills. Whether the risk is retained or ceded, the problem still exists. Because of the indeterminate price of the hedge, the reinsurance price must fluctuate slightly. I am not talking about experience refunding here. My idea would be to outline the pricing methodology and change the price only according to these rules. Obviously changing the price just because the market has crashed is not acceptable. The parameters I have in mind are: (1) age/sex distribution, (2) fund distribution, and (3) a change in the implied volatility of the market—a change in the hedge market's perception about the future volatility of the stock market.

Use hedges to pass market risk to financial markets: a knee-jerk reaction is that buying hedges is risky and the insurance industry should not use them. This would be very

unwise because by attaching GMDBs to your annuity, you have actually already sold a derivative, albeit exercisable only at death. We need to buy hedges on the other side to get rid of this market risk.

Hedging Demands Detailed Information

To adequately hedge this risk, age, sex, and fund allocations need to be updated regularly. Regardless of whether direct writers hedge the market risk themselves, or send it to a reinsurer, accurate information will have to be available and shared. The administration on GMDB risk is going to have to be much more detailed and managed more carefully than life reinsurance. Things do not change very often on a life policy, but any fund transfers made in an annuity are going to impact the hedge, and the reinsurer is going to need to know about it.

A reinsurer can bring to the table (1) a pass on the market risk; (2) the fact that he or she has already worked out the details of the hedge; (3) the fact that he or she has already accepted the mortality risk; and (4) the diversification of fund manager risk—the more varied funds the reinsurer has hedged, the more likely they are to conform to the market expectation.

MR. JOHN SLATER, JR.: I think Tom and Steve did a great job of laying the foundation for my discussion, so I'll try not to duplicate their presentation. A very important point that came out of Steve's discussion is the fact that the death benefit risk is not diversifiable. It does no one in the life insurance industry any good to pass this from life company to reinsurer to retrocession. This risk can be dealt with in another way, and that's really what we're talking about now.

While the concept of hedging has been around for thousands of years, the application to guaranteed minimum death benefits is somewhat novel. The life industry is amassing enormous risk in GMDBs. I'm not sure that we were all aware of risks until Steve brought them up. I'm looking for tremendous change in this area in the next 12 months.

Let's start with the basic definition. What do I mean by a hedge or hedging? Well, I'm not talking about a box bush or primrose in front of someone's garden. I'm talking about the opportunity or the obligation to offset a significant financial liability—to protect oneself from losing by a counterbalancing transaction, as in hedging a bet. Or it is to protect oneself financially by buying or selling commodity futures as a protection against loss due to price fluctuation or by minimizing the risk of a bet. The life industry's unhedged GMDB liability is a $40 billion bet, and the value of financial assets will not decline for any significant period of time. That $40 billion figure is rising at about $10 billion a year.

Now let's look at the characteristics of the GMDB and how they affect risk. Quite simply, the annuitant holds the benefit, and the reinsurer holds the risk. Someone is demanding this feature. Tom suggested that it's not the annuitant. Someone is demanding this feature and the annuitants are getting it. What exactly are they getting?

Let's look at the terms of a typical and somewhat troubling death benefit description. The amount of the death benefit will be the greater (1) the sum of all purchase payments increased by 5% per year, less any amount surrendered (this sounds familiar), and (2) the

contract value. This is great for the annuitants; no matter what happens, they get their money. It's this simple language that leads to the complex problem.

The death benefit description can be written algebraically, as the maximum of payments growing at 5—*P*5 and the value of subaccount assets at death as *D*. Therefore, the reinsurer is at risk whenever *SD* is less than *P*5. The payments are growing at 5%. There are a number of scenarios where the death benefit liability will increase even if the stock market and the bond market remain fairly constant—their price level remains fairly constant through time. And in fact, the situation is further increased or made worse, by the fact that there are a number of expense charges that are deleted from these accounts on an annual basis, and they can be as much as 2%.

To get a handle on the value of the death benefit by itself, I subtract *SD* from both terms. We're left with the max of *P*5 minus *SD* and zero. Now those of you familiar with option pricing, this is exactly the formula for a put option. Put options are one of the many tools in modern finance and at UBS we employ over 100 models to value and price and to trade put options. When subaccount asset values decline, that economic interest goes negative.

News flash! Reinsurers or direct writers have therefore written, sold, or given away put options on a portfolio of subaccount assets, including stocks, bonds, and cash. This may be a surprise to many people at this conference, but perhaps not to the people in this room. That's basically the problem. Now let's consider the solution. How do you hedge the GMDB? First, focus on the GMDB risks. Death benefits come in many flavors: return of premium, interest growth, and ratchet or look back are the most common. Their characteristics must be modeled explicitly. Agree upon the structure of the hedge. Death benefit reinsurers quite often choose to take mortality and lapse risk, and they somewhat unwillingly have to accept investment risk.

As an investment risk manager, I will gladly accept all aspects of the investment risk, which I can pass through to the financial markets—namely, changes in the prices of stocks and bonds. The goals of hedging is to desegregate the total risk into manageable components and then to vest each component with the most efficient handler. As I'm neither an insurer nor a reinsurer, I leave mortality risk and lapse risk to all of you.

Third, structure the hedge based on subaccount characteristics and the realities of the financial markets. Now key subaccount characteristics include the percentage of assets in stocks, the percentage in bonds, and the percentage in cash. Key inputs from the financial markets would be interest rates and implied volatility. Interest rates are quite simply the cost of money and implied volatility is the market cost of hedging a contingent liability.

Fourth, assess the cost benefit impact of hedging after determining the cost of the hedge and the size of the hedge. It is necessary to consider the benefit to the reinsurer. This could be achieved using a Monte Carlo simulation, which simply included the performance characteristics of a put option-based hedging strategy. I say it's simple, because I wouldn't have to do it.

Let's review the necessary inputs in structuring a death benefit hedge. First, subaccount asset amounts determine the size and the type of the hedge, based on the distribution of stocks, bonds, and cash. Second, while mortality is the risk of the reinsurer, the mortality

assumption in the structure of the hedge impacts the size of the hedge—the same is true for lapses. Ages and gender distributions determine the portion of the mortality table that we use, and the death benefit description determines the exact features that must be considered and modeled.

Expenses, insurance, and fund management expenses buy downward the performance of subaccount assets and must be accounted for as well. These inputs allow us to frame out the hedge or what is effectively a series of put options.

Returning to the put option analogy, age impacts the size of the hedge, and the ratio of assets divided by the death benefit impacts the relative cost of the hedge. Second, we assume males hold 50% of assets and females hold 45% of assets. Male and female mortality is applied in proportion to the asset holdings of the male and female annuitants.

We then consider a distribution of assets within subaccounts, which includes 50% stocks, 30% bonds, 15% fixed, and 5% cash. Right off, you can see that 20% of these assets do not need to be hedged. That would be the 15% plus the 5%. The mortality assumption is the key determinant of the size of the hedge. In this example, we will assume 90% of the 75/80 ultimate table. For the sake of brevity, I'm going to skip over a bunch of the remaining inputs. This death benefit, however, is the 5% interest growth variety, the same as Steve just considered, but the simplifying assumption here is that the asset values in the subaccounts are exactly equal to their death benefit level (i.e., where the ratio was 1.0). Remember, if subaccount assets are worth less than the death benefit, then hedging will be relatively more expensive. If subaccount assets are worth more than the level of the death benefit, hedging will be relatively less expensive. Death benefit hedging is a multidimensional problem, but it's very manageable.

Now let's look at the results. The gross cost of hedging, in a theoretical all-stock portfolio, as in the case that we've just described, would be 33.33 basis points per year. The gross cost of hedging a theoretical all-bond portfolio would be 12.5 basis points per year. And I'm assuming here that we're not going to hedge the fixed or the cash portfolio. The net cost of hedging is 8.2 to 14.3 basis points. Now the net cost of hedging adjusts for the existence of a fixed interest option, the presence of cash in any or all of the subaccounts, and for the assumption that correlation of asset returns will be less than perfect. Correlation, incidentally, is one of the least understood and most important parts of this hedging exercise.

Hedging is quite straightforward, but let's consider some of the risks, which fall basically into two categories: structural and those associated with annuitant behavior. I'll touch on only two that have not yet been covered, namely, tracking error and counterparty risk. Tracking error is the risk that the hedge does not exactly offset the liability. This will be a function of subaccount performance and the size of your pool of subaccounts. We have estimated, however, that for most subaccounts this will just be a couple to a few basis points per year. Counterparty risk or credit risk is the risk that a party to a transaction won't be in business at the time that he is called upon to act financially, to make a payment or to receive a payment, especially for long-term situations. It is paramount to choose a counterparty with a strong balance sheet and a high credit rating.

I think the need to hedge and the methodology is pretty clear. Let's consider the question, why does hedging work? First, we analyze in-force details or new business assumptions.

The death benefit liability can be current in the case of in-force business, or it can be prospective in the case of new business. In either case, we designed the death benefit hedge to match and offset the death benefit liability.

Second, we calibrate key parameters of cost in the financial markets, using interest rates and implied volatilities, which ensures that we structure a hedge that can be implemented at a cost that is known in advance. And here's another important conclusion. If you hedge, the cost is known in advance. If you don't hedge, the cost is unknown and potentially enormous.

Next, we translate the aspects of GMDB risk to their option analog, and we employ option-pricing technology whenever possible. At UBS we apply the Black-Scholes or arbitrage-free pricing methodology to the problem of death benefit/risk management, ensuring a hedge with predictable performance characteristics and a known cost.

Finally, we assessed the risks of hedging strategy in advance. Hedging is not without risks, not hedging, however, implies far greater risks. And Tom and Steve have demonstrated these risks very accurately. The results of hedging are greatly reduced variance of returns.

MR. JEFFREY K. DELLINGER: We've heard some fairly exotic solutions to the guaranteed minimum death benefit risk problem posed, and I just wanted to gather the panel's thoughts on why a more simple approach might not work. It seems that this would eliminate the age and sex distribution risk, the fund allocation risk, the portfolio manager underperformance risk, and a few others. If each day we were to just take the guaranteed minimum death benefit, however it's defined in the policy, and subtract off the then current account value (essentially finding the net amount at risk as of that day), and multiply it by the appropriate age and sex mortality rate for that customer, times the net amount at risk, it seems that would work. We've created an explicit deduction from the account value. So anyone familiar with the product understands I've just invented variable universal life insurance on a fairly small scale, but essentially you could even let the customer choose or define whatever minimum death benefit he or she wanted, or have none at all, recognizing that the richer death benefit options would then be assessed higher death benefit charges. What are your thoughts on that?

MR. MITCHELL: That certainly solves the risk problem. Whether that is marketable, I don't know. I would have my doubts about that.

MR. HABEGGER: I'd say the same thing and there's also going to be many administrative complexities associated with that.

MR. MICHAEL GERARD DEKONING: I think you have a real problem with the securities and exchange commission (SEC) people. They won't allow it. I don't believe the SEC would allow that because then it becomes a life insurance contract as opposed to an annuity contract.

MR. SLATER: There's an additional issue, and that is, what do you do with the $250 billion of in-force business that already exists?

MR. DEKONING: I have a question for Mr. Slater. How long was the hedging strategy that you talk about? Was it for the life of the annuity?

MR. SLATER: Generally, discussions have been for 10-, 15-, or 20-year hedges. If you project lapses in mortality forward 20 years, you capture the lion's share.

MR. DEKONING: And in your example that you gave figures based on a ten-year horizon?

MR. SLATER: That was based on the 20-year horizon.

MR. DEKONING: So, the hedge itself renews automatically with no change in price, so that price would be through the whole 20-year period? If you ended up at the end of four years, and had to renew the hedge, and if there was a big risk at that point, wouldn't the hedge cost significantly more than that?

MR. SLATER: Well this certainly depends on who is creating the hedge. The hedge that I envisioned and the one that was used in the pricing would guarantee that cost for the life of the hedge.

MR. MICHAEL W. PADO: You indicated some fairly low numbers for the 5% increasing benefit. But you also suggested that was a net cost. You also suggested there were other risk components in terms of counterparty risk, tracking error risk, and that sort of thing. How does one factor those types of risks back into the cost structure?

MR. SLATER: There are a number of issues here. The net costs have benefited from an assumption of correlation of less than perfect correlation between subaccount assets. I think that's an issue that most people can imagine. I know some of the consulting companies have written papers on this subject. It ends up being extremely important to get from the gross cost to that cost. The subject of costs—hedge costs in general, and the credit quality of hedge providers—brings up the issue of credit in general, and I think that's one that has not been well-addressed in any industry to date. I would advise all of you who deal with reinsurers and with hedgers, and with other insurance companies to consider the impact of credit. If you look at bonds, you'd know that corporate bonds traded a significant spread to treasuries. But I'm not sure that those differences are accounted for in reinsurance or in hedging, but they certainly should be. So the short answer is that it's not really taken into account explicitly, but it should be. But it should be on the side of the people who purchase the hedge or purchase the reinsurance, as much as by the providers.

MR. PADO: I have a question for Steve. As a reinsurer and in terms of that partial withdrawal risk that you were illustrating before, how do you protect yourself in that environment?

MR. HABEGGER: Well, I think there are a couple of ways. One is to try and talk to the ceding company, and to some extent, just the education of the risk could possibly help. A second way is if the reinsurer could base his charge on the guaranteed value and therefore, when they would withdraw like that, that risk would not be passed on to the reinsurer. I don't really like that solution, because all it does is leave that withdrawal risk

with the ceding company. In general, I'd rather keep it on a partnership basis, and have our risks and our interests aligned. Those are my two options.

MR. PADO: As a reinsurer of this benefit, I found that it's rather difficult to leave that risk with the ceding company. With all the hedging inputs that are required—the assumptions as to mortality, persistency, age, sex, fund allocation, and so on, and projection of the put nominal, as it were—it seems to me that it still leaves substantial risk, potentially, with the ceding company for having guessed wrong in the first place.

FROM THE FLOOR: The net cost you gave on hedging on the range was, 8.2–14.3 basis points. Did you incorporate the fact that at the older ages they usually invest in the less risky investments and that at the younger ages, you see a higher concentration of stocks? So isn't your actual distribution weighted by investment and by age?

MR. SLATER: That's a great question, and I'll have to go home and tear apart the model to answer you specifically. There are a number of what I would call second order embellishments, which could be built into the analysis and it can be refined continually.

FROM THE FLOOR: But this dramatically affects the cost we found when we've done our Monte Carlo simulations. But what is that 8.2–14.3 range? Is that by age?

MR. SLATER: No, this is actually taking into account the entire in-force business. Ages range from 30 to 85. The range depends upon the correlation that you assume between the different subaccount asset groups. If the correlation is high, there's a dramatically increased need to hedge the correlation between asset groups used to hedge. And that was just representing sort of typical assumptions of correlation.

FROM THE FLOOR: And there's probably an even greater correlation by age and type of investment that you might build in?

MR. SLATER: I appreciate that point and I'll try to get it in there.

FROM THE FLOOR: Tom, on your distribution of cost where you came up with the $2.40, was that using an average age? Or did that include all the different ages so it would be sort of the worst-case scenario at the highest stage in the worst situation?

MR. MITCHELL: All those figures were based on one cell at age 55.

FROM THE FLOOR: We could have four times that for say, age 75?

MR. MITCHELL: Yes, it would be.