RECORD, Volume 25, No. 1*

Atlanta Spring Meeting
May 24–25, 1999

Session 90PD
Guarantees on Variable Products: How Are Companies Assessing the Risks?

Track: Investment/Product Development
Key Words: Investments, Product Development

Moderator: NANCY M. KENNEALLY
Panelists: BORIS BRIZELI
JEAN-FRANCOIS LEMAY
RISHI KAPUR
MICHAEL SAKOULAS
Recorder: NANCY M. KENNEALLY

Summary: The last two years witnessed a proliferation of guarantees on variable annuities and competing investment products:

- *Guaranteed Minimum Death Benefit*
- *Guaranteed Minimum Income Benefit*
- *Guaranteed Minimum Accumulation Benefit*
- *Mutual fund death benefits*
- *Canadian segregated funds*

To estimate costs of these benefits, companies have moved away from stochastic modeling to capital market pricing approaches. A panel of experts identify the various risks associated with these guarantees, provide an overview of the different pricing methodologies used in determining the costs, discuss whether companies are charging enough for these risks, and forecast the next wave of guarantee designs.

Ms. Nancy M. Kenneally: We have a great panel with us today. Boris Brizeli from Insource Limited, which is a GE Capital Company, specializes in product development. Jean-Francois Lemay is from RGA Financial Products in Toronto. Rishi Kapur, his associate, is also from RGA. Mike Sakoulas is from AXA Re in New York and vice president of Reinsurance Solutions. I am a consultant with Tillinghast-Towers Perrin in New York.

Last year variable annuity (VA) sales reached nearly $100 billion. In today's equity markets there are some consumers who are looking to their nest eggs and thinking that they are willing to forgo some of the high-equity returns to pay for some benefit guarantees in order to protect their investment and give themselves a little

Note: The charts referred to in the text can be found at the end of the manuscript.

peace of mind. By the same token, the VA market is becoming increasingly competitive and insurers are vying for shelf space with distributors who have the option of selling a wide range of different products. Some insurers have turned to innovative product designs and, in particular, benefit guarantees as a means of differentiation.

Guaranteed minimum death benefits (GMDBs) have been offered on VAs for a number of years and have evolved from simply a return of premium to the rich guarantees that we see in today's market, which are usually combinations of an annual ratchet and some type of an accumulation typically at 4–6%. Guaranteed living benefits. Here I'm referring to the guaranteed minimum income benefit (GMIB), which is available upon annuitization, and the guaranteed minimum accumulation benefit GMAB, which are the latest innovations in the VA market.

Equitable was the first to introduce the GMIB benefit. It was actually a combination GMIB/GMDB benefit first introduced in 1996. Since then they've tweaked their benefit a few times, first to unbundle the GMIB from the GMDB and then to reduce the accumulation rate on the benefit from 6% down to 5%. Their benefit has been tremendously successful. We've heard that through their wholesale channel where they're selling through independent agents that upwards of 70% of the contracts sold through that channel elect the GMIB as a rider on their contract.

Over the last year and a half or so other insurers have followed suit. Several have introduced GMIB benefits that are similar to The Equitable's benefit. Northbrook Life, Kemper, Manulife, and SunAmerica are just a few. I also know of several others that are currently under development.

The next benefit to be introduced was the GMAB. One of the first companies to come out with this benefit was Travelers. The GMAB likely has its roots in its Canadian counterpart, segregated funds. Segregated funds are required to provide a return of principal of at least 75% and no more than 100% upon death or maturity after 10 years. Travelers, Providian, Citicorp, and Indianapolis Life all have these benefits out on the street. There are several other companies who are in the process of developing them.

The cost of these benefits to the policyholder are not insignificant and the risks to the insurer are also not insignificant. These benefits introduce a new level of complexity in pricing and modeling.

First, we're going to hear from Boris Brizeli, who's going to discuss some of the risks that these benefits present. Then we're going to hear from Jean-Francois Lemay and Rishi Kapur, who are going to give us the Canadian perspective on these benefits. To give us the U.S. perspective we'll hear from Mike Sakoulas. I'll turn it over to Boris.

Mr. Boris Brizeli: When I say I'm going to talk about minimum return guarantees, risks, and their future, I'm referring to the future of the minimum return guarantees, not the risks. First, I'm going to define the guarantee liability. Then I'll discuss what the guarantee is really worth. And finally, I'm going to discuss the future of the guarantee designs. Let's talk about the main risks. Chart 1 shows the history of the Nikkei index. It's my favorite graph for illustrating the risks because it really amplifies them. Here we see that somewhere in the early 1990s the Nikkei was around 40,000 but by the end of 1992 it was pretty much cut in half and now is hovering around 16,000. It has recently reached about 17,000. What I'm displaying is not the actual payouts, but what the exposure to a particular risk is, so if you have a product that resets or ratchets at the peak around 40,000, then today all the amount that you would have at risk is all that distance at the top right. If people die where you see the arrows drop down, that's the amount you're going to pay out. If you have an accumulation-type benefit that's paid out after, let's say, ten years, then what's indicated by the market drop at the end of ten years is what you're going to pay out.

Let's talk about the different types of guarantees. The first type of guarantee is the fund basis. Essentially the guarantee is the sum of the individual subaccount guarantees. Transfers are easily modeled in such a setting. The guarantee in subaccount (i) decreases if you have a transfer and the guarantee in subaccount (j) increases. That's how transfers work in this guarantee. Whether you do FIFO or LIFO, that can impact the fund basis somewhat differently around the maturity date.

The second type of liability is the account basis. Here the guarantee is on the sum of the individual subaccounts. Not all funds move in the same direction even though many believe there is very high correlation in the markets today. Funds can move in different directions, and you're offering a guarantee on the sum of the funds. The guaranteed amount will always stay the same in this case and the maturity date can stay the same depending on whether you use an FIFO or LIFO basis. Let me refine the maturity date definition somewhat. There are two ways you can define it. One is to define it by issue date such that maturity will be measured from issue date so every deposit that goes into the contract gets a shorter and shorter maturity date. With a deposit one year from the start of the contract, essentially the tenure of the put will be one year shorter. What you effectively get in a guarantee like this is a combination of a limited amount of puts of duration from one day to ten years. There are some really interesting arbitrage possibilities with products like this; they actually exist in Canada and are still being sold. It's really an interesting product in that it presents true opportunities for which you can arbitrage the guarantee against the market. Another way to define the maturity date of a guarantee is by the deposit date; that is, every deposit gets its own maturity date. That's a more common type of design. In a Canadian product such as a segregated fund a typical period is ten years, so every deposit will get a ten-year tenure.

Some more refinements here. Management fees and guarantee charges impact the cost of the guarantees because they're essentially a negative drag in your stochastic process so they increase the cost of the guarantee. Essentially this is equivalent to an additional negative return. There are different ways of modeling that. I used to believe it's best to model it as an "in-the-money" put, but I've changed my mind. There are better ways of doing it, and I think Rishi will talk about those next. Let's talk about the reset option a little bit. Here in the U.S. you have the ratchet-type guarantee. In Canada we have resets where you can elect to essentially freeze the new market level. The basic idea is the same whether they're elective or not in that what you have is a cost that's not dissimilar to unamortized expenses. That is, you pay for something with periodic fees and you're allowed to exchange that and start anew. You can leave behind a positive value or a negative value of the guarantee. So, essentially, when you reset a guarantee level it's as if you're taking the money out and putting it back in. That's effectively equivalent to leaving something unamortized behind because you're paying for it with a periodic charge. So it's equivalent to a put back to the insurer of the book value of the put. Now, depending on the type of guarantee that exists what happens after the reset will be different, but not only that, even if the price after the reset is not guaranteed you still need to hedge or somehow address this risk. It doesn't matter that after the reset you can change the fees. It is still important to hedge this option and a lot of people believe that if you have the freedom to reset the price after the reset that that's meaningful; not necessarily, you still need to hedge this cost.

Let me talk about some liability risks. I've seen this topic discussed so many times in SOA sessions and CIA sessions, so I'm going to go through this rather quickly and assume people are familiar with what I'm talking about. What is market risk? I'm going to define it the following way. It's the exposure arising from the trigger of the guarantee that is set off by the performance of financial markets. In a segregated fund context the market risk can be a result of several events: reset, lapse, death, or maturity. Maturity risk depends on the level of existing guarantees, the time to maturity of the guarantee, and the current asset value. These are all the usual option pricing parameters. However, this is where a key portion of any good modeling comes in. If you're modeling a type-two guarantee where the guarantee is on the sum of funds, the maturity risk is a function of the correlation between subaccounts comprising the assets. It's a function of net amount at risk. Rational lapses are essentially resets because they're resets that are exercised optimally. I should be very careful here. A rational lapse is a lapse that says a person will only take away his or her money from the fund and redeposit it when it is most optimal for him or her to do so. Maturity risk is also a function of interest rates, currency, etc.

Basis Risk

What is basis risk? The definition of basis risk is exposure arising from the hedging of the guarantee using an asset that is different from the asset supporting the contingent claim. Essentially, if you assume that there is a degree of correlation between two different assets and you calculate your risk-management parameters on an asset that is different from the one that's being hedged, you have basis risk because the correlation changes dramatically over time. In any fund that has an index there would typically be no basis risk. However, if you do not have a market index that you're hedging you have a basis risk. Your return distribution is composed of two components, and we assume that any given actively managed fund tracks a particular benchmark. So there are two types of returns. There's the return that tracks the index and there is the incremental return that can be positive or negative from the active fund management. Now that particular incremental return has a distribution, and it's that distribution that you need to model to know what the basis risk is.

Mortality Risk

It's when the fund is underwater compared to the initial levels and somebody dies. That's where you pay off—it's a top-off. One of the key issues here is antiselection. Why? Because there are cross-subsidies from the way you're pricing for it. You're charging one fee to all ages, so antiselection is quite possible. There is significant evidence of this antiselection by now. In fact, there are several advertisements in the distribution system and in newspapers telling you to buy VAs or segregated funds at older ages. You must be able to not only estimate the appropriate level of mortality, but also the distribution of risk.

Lapse Risk

Lapse risk is really not that simple because I define two different types of lapses as rational and irrational. On one hand, you can have lapses that release mortality and maturity liability— those are the good ones—but for that you have to forgo fees. That's what rational lapses allude to. Lapses depend on the kind of scenario you're in. Lapses are extremely path-dependent and depending on the path that you experience you get very different lapses. It's no easy task to model how lapses work. For example, a distribution system is an extremely key component in modeling lapses. I've seen some models of policyholder behavior, and a lot of people think that you can apply a universal model with a few parameters saying, "OK if the put is in the money, people will lapse less; if it's not people it will lapse more." It's not as simple as that. There are a lot of irrational things that people do. In fact, lapses can be much higher when the put is in the money because people are scared. So you may in fact expect higher lapses. And that's not at all an unreasonable expectation in specific distribution systems, so be very careful there.

Lapses depend on several other things. What is the legal form of the fund? Is it a registered fund or not? Reset level, age, and risk profile are other variables. Irrational lapses, which are the best lapses, are really good for you because what they can do is cut the maturity guarantee in half because those people lapse suboptimally. It's really a good thing.

The key point I want to stress about liability risk, specifically policyholder behavior, is this: You can only manage these things; you cannot really control them. Some risks you can control through how you market and design the product. I'll give a few examples of how you can manage some of these guarantees, but what you're dealing with is behavior. You're dealing with humans and that is one of the toughest things to deal with. But the key point is this—you cannot hedge these things. You must make assumptions, and for any assumption that you cannot hedge essentially you'll be wrong one way. You'll very seldom hit it, and you will either be wrong on the upside or the downside, but every time you make an assumption that cannot be hedged that's what you're faced with. We're very used to that as actuaries. We make long-term assumptions all the time. We know we're not going to hit them exactly, so we have several tools for dealing with that. Well, we have this very same thing present here. These kinds of things simply cannot be hedged.

I'll talk a bit more about reset risk. The really neat thing about reset risk is that it works in the other direction of the basic risk. You have an option in this product that's in essence the best of both worlds. If things go down you get a payout; if things go up you get to reset. The resets were designed (they were invented of course in the U.S.) to keep the guarantee in the money. People wanted these contracts to persist. Well, I don't think it was necessarily a good thing in retrospect.

The way to deal with reset risk is to define what optimal behavior looks like and then hedge it. If you're going to be using inefficient policyholder behavior in dealing with resets specifically, you must make sure that inefficiency is an explicit assumption in your model. This way you know how much risk you're taking on in relation to what the optimal behavior can be. Reset risk is very scenario-dependent. For example, if you have a year where you have a lot of resets, it's quite likely that next year you're going to have a significant level of reset activity. On the other hand, if you had a year that was pretty flat and you didn't get a lot of resets and next year it increases, it may well be that there is an apathy factor. If people didn't reset last year, they're not as likely to reset next year even if there is an opportunity to reset the risk. Now I don't know if that's necessarily a true model; however, it certainly is true in some distribution systems. In others you have extreme examples of antiselection. You have agents who have a pretty decent idea around when they should reset and every so often they send a list of policyholder names to the company saying I want all these people to reset. It really depends on how you distribute this stuff. In more naive settings—for example if it's an over-the-counter retail product distributed in very low amounts to, let's say, bank customers—that kind of product would expect very naive behavior. That isn't too sophisticated. And, as I said before, it moves in the opposite direction of the basic risk.

Fund Distribution Risk

That's a risk where you can have poor market timing. Let's say the equity market dies and people transfer everything into something very low-yielding like a money market fund. That's really a good example of poor timing because in theory, you should just stick around and let the market catch up with itself, but by putting the money into the money market what you are in fact doing is locking in a loss, so that's problematic. Usually policyholders don't diversify enough. Another example of fund distribution risk is that the very presence of the guarantee encourages risky behavior. Time and time again I've heard agents say the following things: "I have a substandard risk on my hands and I'm going to put him into the Asian Fund." Why? Because there are guarantees there. What's the worst-case scenario? That this substandard will die and get his money back. That's the worst-case scenario. If he gets a lot of upside, well, great. In this type of fund you have substandard risks with a very real benefit.

Fund allocation determines price. If you make some cross-subsidy assumptions, that has a big impact on the price. You cannot really hedge a fund allocation; that's a behavioral issue.

Fund Management and Active Management Risk

The Financial Post published a really neat report. It said over a ten-year period, two-thirds of active fund managers in Canada will do worse than the index they're trying to follow. This gives you an idea about how the incremental return that I spoke about before is distributed. In the U.S. the numbers are even worse. This is not at all to say that U.S. fund managers are worse; I'll let you make that conclusion on your own. However, I think this illustrates to you that the distribution of incremental return can have a very significant impact on the guarantee value. On top of this we have the management fees and guarantee charges drag. The guarantee may alter investment allocation. What that results in is the volatility of the guarantee increases, so you can have a situation where the bogey performance can end up with a put. If you bought an "at-the-money" put it would actually not pay off, but the fund itself would pay off because of the guarantee drag.

What is the relative magnitude of these different risks? For example, I have market risk as the biggest risk if it's left unmanaged; however, if you manage the market risk it's the one that's most susceptible to management. You can significantly decrease the value of that risk. It's a huge risk if you leave it unmanaged; it's a catastrophe risk—high severity, low frequency. If you apply risk-management techniques to it, that's the one over which you have the most control. But, for example, look at another risk at the other end of the spectrum—it's a lapse risk. It's not as significant as the market risk, but it's not very susceptible to risk management. What I'm trying to illustrate is the relative magnitude of the different risks and how susceptible they are to different types of risk management.

Other types of risks include model risk. Model risk is really worth mentioning just to say that what we get in reality is not what our modeling tells us will happen. A lot of model imperfections have a real cost associated with them. Whatever model you use—you can use actuarial-type models or you can use models that look at capital market behavior—both have significant model risk. For example, for capital-market-type models you have liquidity risk, called liquidity holes. You have absence of fair prices where the market gaps on you. The bottom line is you may get the process right, but not the parameters. So even the very assumption of the process may not be a good one. There's correlation risk and large event risk; that's where the market dives 500 points.

Counterparty Risks

Those are present if what you're using as your risk-management approach is reinsurance or you're using over-the-counter derivatives. You should pay particular attention to counterparty risk. Why? Because that's the risk that if things go wrong they go really wrong. They don't go wrong a little.

What are these guarantees really worth? The clear answer is, it depends. Why does it depend? I'll tell you my idea about this. There's no universally accepted methodology for pricing. There are a lot of people who believe they have the right methodology, but there's nothing that's universally accepted. I'm not saying that universal acceptance is what's needed. I'm saying it's something that's really useful if you want to answer that question. There's no universal view on how to manage this stuff, so again that prevents me from answering that question. One other thing that prevents you from answering something like this is the following fact: If I go to the capital markets and I want to see what the price of IBM stock is, it's directly observable. You cannot observe the key market parameters for this long volatility. It is simply not observable. Not only that, but there is no market being made in it. Nobody's buying and selling long puts, which is what most guarantees are. In absence of a market it's very tough to say what something is worth. When you look at the fair value of insurance liabilities it's very tough to put a fair value on it. This is because there is no active bid and ask. Another unknown is that we have no idea how customers will behave. Even if we have a decent, let's say, 20 years of history of how the market and policyholders behave, is that really a lot of valuable data? One must be very careful from generalizing from what is essentially one single scenario. Even if you have 1 scenario that runs for 20 years it doesn't tell you a lot about how people will behave in the future. At best you can have some really good ideas about extreme types of behavior, but a lot of local types of behavior will be influenced by factors that you may not even know about to date. That's a tough one.

I'll lighten up the picture a little bit. I'll talk about some future design trends and just give you some general directions where those guarantees are likely to go. One is mass customization. That's a natural direction for any of these products, GMIB, GMAB, or GMDB, to go in. I believe companies will start developing a menu approach. Where do you want your maturity risks? How many resets per year would you like? They'll define several parameters and let you choose, and there will be a price assigned to every component. That's a natural extension of what I think we're seeing in the market today. This is essentially a predication that says we'll see a little bit more of what we're seeing today in a little bit better package.

Product line extension is a really interesting trend. We're already seeing it because we went from accumulation products to payout products, but I believe there is a lot more room for innovation in terms of product line extension. A very exciting arena, of course, is universal life (UL). We're already seeing several examples of minimum return guarantees applied to UL in Canada and the U.S.

You can extend the different decrements for your design. As a very simple example, instead of death use critical illness. You'll top off not in case of death, but in case of critical illness. You need some creativity about designing that one because it's a living benefit. It's not a death benefit, but it certainly presents some really neat possibilities and new exotic derivative packages. The basic drivers of these guarantees are people's needs: What kind of insurance they need, how they want to allocate their investments, and how you can arbitrage the perceived value that people give to a particular type of guarantee for it's actual value. If the perceived value is significantly higher than the actual value, you have a market. And in designing these kinds of guarantees that's what you need to concentrate on analyzing.

Mr. Rishi Kapur: I'm from RGA Financial Products and, together with my colleague Jean-Francois, we're going to talk a little bit about the Canadian context in all of this. The Canadian products are somewhat more complicated than the U.S. products. We're going to talk about the products and their pricing. I'm going to touch on actuarial models versus capital market pricing models and Canadian segregated fund guarantees. My colleague Jean-Francois is going to get into some of the pricing methodology and risk management for these products.

The two broad categories that we've divided pricing models into are actuarial and capital markets. On the actuarial side these models are based on historical analysis of data regarding mortality, morbidity, policyholder behavior, and investment returns. As actuaries, we're quite familiar with these types of models. The benefit prices that are developed by these models are based on expected claim costs and provisions added for risk charges, adverse deviations, and things like profit margins. The capital market side takes a different approach. The capital markets pricing models are generally based on current market information about stock prices, interest rates, and volatility observed in the market every day. The benefit prices are based on expected hedging costs and not on expected claim costs. If you somehow see hedging as production, then the prices of these products are based on how much it costs to produce these products in the financial markets and not actually what you expect to pay out. This is somewhat of a different way of thinking from the actuarial approach. The market price of risk is implicitly included in the pricing methodology for these products. In the actuarial side, explicit provisions are added for profit margins and adverse deviation. In the capital markets it's implicitly included in the pricing.

What are the advantages and disadvantages of these methods? For the actuarial side the advantage is that it is the best estimate for expected claim costs if the risk is unhedged. If you are holding the risk naked, or as some people like to call it, self-insuring, you're not hedging, and hedging cost is irrelevant for your model. Actuarial models work well for econometric analysis and projections, especially for things like dynamic solvency testing. The actuarial approach is often the only way available if risk cannot be hedged—risks like lapse risk and mortality risk. What are the disadvantages? In general more inputs need to be estimated. This requires a view on expected returns on stock. By this I just don't mean stocks, I mean things like interest rates and currencies as well. The actuarial models do not reflect the costs of production (hedging). They also create arbitrage opportunities in the sense that the way the products are priced you can actually do a lot of significant underpricing compared to capital markets using these actuarial models. What about the capital markets models? Well, they do reflect the cost of hedging. They can be used to develop and execute hedging strategies for risk management unlike actuarial models, which cannot be used directly for risk-management purposes.

Some of the disadvantages of the capital markets models are that the price is not reflective of the cost if no hedging is done. Again, there is a distinction between hedging and not hedging. If risk cannot be hedged, the capital market pricing models may not work. They do not work for lapse or mortality. The calibrated parameters are more volatile and subject to supply and demand in the capital markets. Again, it's a lot more volatile than on the insurance side where you may value your liability once every three months or once a year in some cases.

What is a segregated fund product? Before we actually get into the pricing of it, let's look at some product features. The maturity is ten years from the deposit date. The deposits are guaranteed on both death and maturity, not like in GMDB where the guarantee only kicks in when somebody dies. The guarantee is calculated on the portfolio with this unique feature called a reset not available on current U.S. products. The account value can be reset at most twice a year and every time you reset the current account value it becomes the new guarantee. So it's like a ratchet, but the maturity is now pushed back ten years from the reset date. Every time you reset you get a higher guarantee, presumably, and the maturity date gets pushed back. That becomes difficult to model. Now, looking at the stock market returns in the 1980s and the 1990s, it's often felt that if equity markets do as well as they have been doing, the risk should be very low in these products. A 5% or 6% roll-up when markets are going at 15% isn't really a big deal. We decided not to focus on the 1980s and the 1990s. These days a lot of people are talking about how the economic environment is similar to the 1960s with interest rates low and inflation low. What happened in the 1960s and 1970s? Chart 2 shows the monthly close of the Dow Jones Industrial Average from December 1958 to December 1982, which is about 24 years, a fairly long period of time. As you can see, the market didn't really do very much over this time. It oscillated up and down. There were 5 occasions during this period that the market dropped more than 20%, sometimes dropping as much as 40%. Every time the market moves up, people would reset or the ratchets would kick in, but then the market goes down and people die and they collect the money. Oscillating up and down can be very dangerous because every time the market goes up, the guarantee amount goes up, but when the market goes down you have to pay out money. This is just to give you an idea of the kinds of risks that are out there. Can this happen again? We hope not, but nobody wants to bet the company on an optimistic scenario.

I'm going to talk a little bit about the capital markets pricing approach. These are the ones that we use in our work. One of the key components of that is something called a growth rate of mutual funds. What do you expect the market to do over the next 10, 20, or 30 years? If you're hedging the risk it doesn't matter what the market does. The expected hedging cost, which is really the key for you if you're hedging, is obtained using the risk-free rate of return as the growth rate of the asset—any asset, whether it be a bond fund, a balanced fund, an equity fund. You assume that they all collect the risk-free rate of return. It's not very intuitive, but this is true no matter what the actual growth rate is, and this is dictated by arbitrage-free pricing. Right off the bat if you start with this process you don't have to worry about how the markets will perform. Whether the markets are going to go up 5%, 10%, minus 5%, or minus 10%, it doesn't matter.

One of the most common assumptions used in the capital markets is something called a log-normal assumption. Capital markets use this assumption for the distribution of future stock prices. It's more of a convention than an actual estimate of future stock prices. Nobody really believes that markets behave in a log-normal manner. However, this simplified assumption can be modified to take into account things like volatility smiles and skews observed in the capital markets. The capital markets have a way of using a log-normal assumption and modifying it slightly to take into account things like stock tails. Jean-Francois is now going to talk about the pricing methodology and some of the risk management that we do with these products.

Mr. Jean-Francois Lemay: Rishi talked about the main points of our capital markets models approach to pricing with the growth rate of the fund and the log-normal assumption. To get a price for these products, you would just generate future market scenarios using a Monte Carlo simulation. You would do your standard simulation using the log-normal assumption and the risk-free rate of return as a growth rate. You would generate these to correlate together. You can't just say I have an equity fund and a bond fund and generate a bunch of these and a bunch of those and add them up. You need to generate them to correlate together so there's an additional assumption that you need, the volatility, and the correlation between funds. So now I have all my funds generated into the future. I also need to look at policyholder behavior. Boris talked about this earlier. What is rational policyholder behavior? That would be the person who exercised the reset option when most optimal. In order to know when is it optimal to reset you need to be able to judge at which point it is worthwhile to have a longer term guarantee. If I reset, I push back my maturity benefit, but I get a higher amount. I lock in a gain, but I push back my maturity benefit. I need to be able to model this so I know what is basically the worst policyholder behavior from the insurance company perspective. This can also be applied to lapses. If you have a GMDB product that doesn't have a reset feature, lapses can be viewed as resets and selected lapse. If people would start to act rationally and lapse when it's worthwhile for them to drop out of the product then you would have to be able to model that. That's the second step you need to do once you have all these scenarios generated. Once you have that, you can have a reset boundary. Now all I need to do is run through and take the average cost of all these scenarios. In practical terms, when you're pricing you don't want to just add on your pure expected hedging cost, you want to add some margins to cover your capital costs, transaction costs, misestimation of assumptions, basis risk, and so forth. These depend on your risk tolerance and, obviously, what the market will bear.

Once you've priced it and sold it, what can you do for risk management? You have four basic options. You could run it naked or self insure it. Obviously there are no costs or couterparty risk associated with that, but you're totally exposed to a risk that could prove to be very expensive. You could purchase reinsurance, if available. I've seen, especially in Canada, that reinsurance may be difficult to get, but that would completely eliminate your risk. It could be expensive depending on the reinsurance quote that you get. You do get a counterparty risk, as Boris mentioned earlier. You must look at counterparty risk because if there's a big market downturn the reinsurer will have to pay up on all its claims at the same time. It's important to look at that. It could be difficult to unwind if you want to get away from the reinsurance, and it could affect the viability of the product. That is, if your reinsurer pulls out, you may have to pull the product if you're not able to manage the risk yourself or are unwilling to take the position naked. You could go out and purchase put options to eliminate some of the risk. It's not perfect risk elimination and could prove to be quite expensive since you basically pay an investment bank to hedge themselves on the put option you're buying. You could over- or underhedge, and you will need to rebalance depending on how many lapses or how much new business you get that would trigger a bid-offer spread on these things. These options are not liquid at all. In fact, it could prove difficult to unwind. If you do dynamic hedging, that could prove to have moderate costs, and the risk is managed. You don't have a lot of counterparty risk. If you're doing dynamic hedging you would go to exchange-traded options or futures, which have very limited counterparty risk. You stay in liquid instruments since it's easy to unwind, but the risk is not eliminated—it's managed, but you still retain some risk on your books. Dynamic hedging is a combination of your inputs, your model, and your hedges. Any of these can be a source of risk, and you have to get everything right for your dynamic hedging to work.

Insert 1 shows a comparison between dynamic hedging and asset/liability matching just to show that duration is the equivalent in dynamic hedging of delta and convexity gamma.

INSERT 1
ANALOGY BETWEEN TRADITIONAL ALM & OPTION RISK MANAGEMENT

Bonds	Options
• Duration: $\frac{^{a}\text{Bond Price}}{^{a}\text{interest rate}}$	• Delta: $\frac{^{a}\text{Option Price}}{^{a}\text{underlying value}}$
• Convexity: $\frac{^{a}\text{Duration}}{^{a}\text{interest rate}}$	• Gamma: $\frac{^{a}\text{Delta}}{^{a}\text{underlying value}}$
	• Vega: $\frac{^{a}\text{Option Price}}{^{a}\text{volatility}}$

If you look at Chart 3, I'm not hedging. This is my distribution basically of what are the possible scenarios. If you look at the dotted line, I have a very high probability that I will get little or no losses when I'm not doing anything, but there are very long catastrophic tails. There are chances that I will get a very severe loss. If I'm hedging I can get my distribution a lot tighter. Managing the risk will basically cut the very catastrophic tails. I'm still not perfect. I have some remaining distribution that could be expensive, but I've really cut this down. The darker line shows what's expected if I'm managing. This is a more realistic scenario for segregated fund guarantees in Canada and most of the funds in the U.S. where you do have basis risk, meaning that you're going to be hedging the guarantee based on actively managed funds so I can't go and buy an option directly on my fund. I have a difference between what I'm going to be using for hedging, be it a future or a put, and any short positions I'm netting in that actively managed fund that is going to keep changing. So I don't have a perfect fit. Here's just an example to show that you do get a wider distribution so your hedging is not perfect when you go into dynamic hedging. You need to know that your results are not 100% guaranteed; but even in this case you don't have any catastrophic tails, so you've cut down on your worst-case scenarios. The risk may not be eliminated, but it is managed.

Reserving. I want to look at the 95th percentile of my distribution to get an idea of cost. Without hedging, it would be a 5.2. If I'm hedging, since there's no basis risk, that number would only be 1.2, and if I'm hedging with basis risk it would be 2. So you can see that the catastrophic end has been greatly reduced. And if you're reserving or putting up capital requirements on a percentile basis, you would have a considerable savings from dynamic hedging.

We've talked a lot about dynamic hedging. But when you're actually doing dynamic hedging how does that work? You start with a model and all kinds of input. Many inputs will come just from the in-force block, for example, lapses and mortality assumptions. Then you need additional inputs that come from the market such as the interest rate, volatility, and correlation. These, which are very important parameters, need to be estimated. Once you get that your model will give you the risk and the risk sensitivities delta, gamma, etc.

That's all fine and dandy, but it still doesn't tell me what do I need to buy or sell because I have an actively managed fund. I need to go through a fund replication routine, meaning that I have all these actively managed funds and I'm going to have to go and figure out which assets I'm going to use to do my hedging. I need to find assets that are liquid enough but follow closely enough with that underlying fund so that I can use that to do my hedging. I would create my hedges and, between my fund replication and the risks that come out of my model, I would do my rebalancing. That's really when you can actually do your transaction and your dynamic hedging.

Mr. Michael Sakoulas: My discussion is going to focus on the GMIB rider that's attached to VAs. I'm going to compare a simplified capital market type of pricing to what we do at AXA Re, which is, I guess, somewhere in between a true capital market price and an actuarial appraisal method. The first thing I'm going to do is to define the GMIB rider and then discuss what sort of risks the insurance company is bringing on the books by writing such guarantees. Then I'll discuss the inherent risks of the GMIB and Black-Scholes to evaluate the price of a GMIB under a sort of simplified framework. And, finally, we'll go over how we do it at AXA Re.

Let's define the rider. Remember, we're talking about an income-benefit guarantee during the accumulation phase. The insurance company is essentially setting it up so you can lock in your annuitization basis during the accumulation phase. Here you have your minimum annuity purchase rates. These have always existed in life insurance settlement options. You have your conservative mortality and your low interest rate (about 3% or so). These benefits were rarely "in-the-money" unless you had a really old block of business. But the new wrinkle here is the guarantee on market performance. This is what we call the income-benefit base. Here you can either guarantee the interest on funds or you can have a ratchet benefit such as maximum anniversary. You can even do combinations. Now what's happening here is the annuitants are comparing what they can get with the income-benefit base and the minimum annuity purchase rates to what you get with the account value and the current annuity purchase rates. They then decide whether to annuitize or not. It's not an automatically exercising option; it's up to the annuitant to elect. And it's not always purely a financial decision. It's here where we can take advantage of the inefficient exercise. But election can only occur after a waiting period. It's usually seven to ten years, or it could even be longer. From a product-design perspective, I think it makes sense to have multiple waiting periods for your product just for the simple fact that the longer the waiting period the cheaper it is. So you could offer the longer waiting period to your younger contract holders who are not even thinking about annuitization for a long time.

The other thing is they could only elect within a window of each anniversary after the waiting period. Here is where insurance companies are trying to limit the availability because there is a potential for excess utilization. And, finally, there could be some attained age requirements. Here you have the GMIB. It's an exotic option and it can only be exercised by annuitizing in some sort of life contingent annuity. When that happens the contract holder loses liquidity and actually loses control of his or her money, or maybe more importantly the agent loses control of their money.

Anyway, let's look at what's at risk. The risk amount is the income-benefit base times the ratio of the minimum annuity purchase rates to the current annuity purchase rate minus the account value. Now the great thing here is that the minimum annuity purchase rates are pretty conservative, so that ratio runs about 70–75% most of the time. It actually relieves some of the pressure of the underlying interest rate guarantee on the income-benefit base. In terms of claim cost it's dependent upon duration, age, and the annuitization rate. There's no claim cost before the waiting period or until the attained age requirements are met, but after that it's equal to the risk amount times the annuitization rate. This is a pretty risky product not only because of the investment guarantee, but because of the annuitization. And since there are no empirical data on such annuitizations it makes some people nervous.

These two risks are the main ones (investment and annuitization), but there are other risks that we have to deal with. We categorize them as follows: market-related risks, behavioral risks, distribution risks (risks created because of the business profile that you bring in on your books), and other risks. Let's look at the market-related risk. Obviously we have market return since we are guaranteeing market performance upon annuitization.

Volatility

These are options so their quotes vary directly with the level of volatility. There's basis risk since most of these funds are invested in a basket of funds, and these are usually not perfectly correlated with the underlying index of the hedge. Interest rate risks have an immediate impact on option costs. Revenue risks. You have to finance the up-front costs of the options or the hedge that you're setting up. You're doing this over time while the charge to the contract holder is annual. In poor performing scenarios there's a probability that your actual revenue starts to decrease. The next risk is fund manager risk. A typical VA product has between 20 and 40 different funds that are offered, so poor performance of a single manager, whether it's because of error or maybe some wicked behavior on his or her part, can have a material impact on the overall performance. However, this risk can be diversified from the reinsurer's perspective.

Let's look at the behavior risks. Obviously the annuitization rate is a major behavior risk.

Mortality Risk

Actually it's longevity risk because it increases both persistency and the value of the guarantee since it decreases your current annuity purchase rate levels. That's the denominator in that ratio that we spoke about before.

Persistency Risk

The longer people stick around the greater the potential for future claims. Revenue risk also appears here and this is due actually to early lapsation before the waiting period, so you're missing out on that revenue. Next we have the take-up rate or the portion electing the GMIB rider. Higher than anticipated take-up rate actually increases your risk exposure.

Transfer Activity

This is mostly an irrational behavior. An example is people dumping their money into money market or fixed-account funds after a market drop. Finally, spousal continuance. Some insurers allow the spouse to continue the contract upon death of the contract holder, so instead of the risk ending it could continue. Let's look at the distribution risks now.

Asset Allocation

Both riskier fund allocations and high allocations to money market or fixed funds can be problematic.

Contract Size

Here there's a lot of fixed contractual fees in dollars. It can have a material impact on small contract sizes. You would hate, from the perspective of large contract sizes, a $3 million contract to annuitize versus the one that's $10,000.

Age

This is mainly a concentration of business at certain key ages, so you may have contract holders behaving similarly. Single life versus joint. This is a similar issue to spousal continuance. And, finally, time. This is another concentration risk. The key example here is if you offer the GMIB rider to your in force. You shouldn't allow them all to elect the benefit or start them off at the same day. You should start them off, for example, on their anniversary. You want a distribution of your business over time. This is what we call a strike diversification. The other risks include the data. It's important to monitor the risks, so timely and accurate data are important.

Modeling Error

As with any pricing, you're always concerned about the integrity of your model, the tax law changes, and the regulatory environment. Any sort of change that could create incentive to annuitize is definitely a risk. So we've defined the guarantee and we've also discussed some of the key risks. Let's price one of them using Black-Scholes.

We're going to look at a simplified deal so you can get a taste of what's going on. The first thing we'll do is calculate the strike prices. The great thing about the strike prices is that they're completely dependent upon the product design, so whether you're an actuary or an investment banker you should come out with the same strike prices. However, we definitely differ on the notional amounts. Here it's a little bit more subjective because it depends upon persistency and annuitization or how the option is elected when it's in the money. Next we'll calculate the up-front costs using Black-Scholes, and then we'll convert it to an annual charge. Finally, we'll look at the price sensitivity as we make changes to the underlying product design.

Let's set up the problem now. We have the GMIB with a 5% roll-up capped at 200% of net considerations. It has a 10-year wait, and the ratio minimum to current rate is held constant at 80% for all years. There are total fees of 250 basis points. This also includes the cost of the option. Funds are invested in Standard & Poor's (S&P) only. The annuitization rate is a flat 10% and whether it's a penny in the money or that much in the money, only 10% of the people are going to elect it. Lapse and mortality are held constant at 6% in all years. Our capital market assumptions will assume a flat yield curve at 5.5%, and a flat volatility curve at 29%.

Let's set up the strike prices. As I said before, they're solely dependent upon the product design. When you take a step back, think about what has to happen. The market has to beat the roll-up percentage and the fees, so we're talking 750 basis points a year. But you have to remember that the ratio of minimum to current rates was about 80%. That actually will relieve some of the pressure of the interest rate guarantee. What you do here is look at the net amount of risk formula, and you solve for the return on the market that equates the account value to the income-benefit base times the ratio. What you actually get are the equations in Insert 2 when the income benefit base is less than twice net considerations (since it's a 5% roll up this is the strike price for the first 14 years). Then after that, when the cap hits, you have that equation or just combine the two and get that funky little equation. Now to look at it graphically in Chart 4, the Xaxis is time and the Y-axis is the strike price in terms of percent. Notice the rally at 15. The slope of the line actually declines a bit and that's just due to the cap of the income-benefit base. At that time the market only has to earn fees and you're set.

INSERT 2
GMIB COST USING BLACK-SCHOLES
STRIKE PRICE IN YEAR *n*

$$= \left[1.05 * \left(\frac{MAPR}{CAPR}\right)^{\frac{1}{n}} + f\right]^n \quad \text{when IBB<2*Net Considerations}$$

$$= \left[\left(2\frac{MAPR}{CAPR}\right)^{\frac{1}{n}} + f\right]^n \quad \text{when IBB=2*Net Considerations}$$

or combine

$$= \left[\left(\min\left[1.05^n, 2\right] * \frac{MAPR}{CAPR}\right)^{\frac{1}{n}} + f\right]^n \quad \text{for all } n$$

Notional amounts are a little different. We have to actually generate them stochastically. We know they're zero before the waiting period expires, but after that we have to simulate the equation shown in Insert 3. Essentially this is just a ratio of GMIB claims to the payoff of a put option with one dollar notional. You actually simulate that for all time after the waiting period and take the average cost of all scenarios. The great thing here is that regardless of your capital market assumption that ratio is actually pretty stable. When you graph out your notional amounts, you actually get what is shown in Chart 5. It kind of makes sense. You have people dying, lapsing, and annuitizing, so you need less notional as time goes by.

INSERT 3
GMIB COST USING BLACK-SCHOLES
DETERMINE NOTIONAL AMOUNTS VIA STOCHASTIC PROCESS

=0, when t#waiting period

$$= Average_l \left[\frac{Max\left[IBB_t * \frac{MAPR}{CAPR} - AV_t, 0\right] * Ann.Rate_t}{Max\left[Strike_t - \prod_{j=1}^{t}\left(1+i_j\right), 0\right]}\right] \quad \text{,when t>waiting period}$$

Note that notional values are relatively insensitive to changes in market return and volatility assumptions

When you put all those things in the formula, you come out with 9% of premium or 60 basis points on average. You might be thinking to yourself, wait a second everybody's charging 30–40 basis points, where did you get the 60? In this example, all the funds are invested in the S&P. In reality, you have about 20% in money market and fixed income and that helps drive down the cost. You also get the benefit of diversification, so actually the notional amounts and the strike prices that we generated here are a bit on the high side. Also, the ratio of the minimum rates to the current rates was running about 80% in this example. In reality, it runs a little bit lower so that actually reduces both your notional and your strike prices.

The next reason was that the lapse and mortality was set at 6%. As your business ages and matures that rate could get much higher. And finally, we assume that annuitization occurred as soon as the benefit was in the money. In reality, customers need some sort of incentive to give up liquidity and control of their money so they're not just going to elect it when it's just in the money.

Let's look at the price sensitivity as we change some of the underlying product features. We had a 10-year waiting period. What happens when we have a 7-year waiting period? Our price goes up about 20%. That makes sense since our claims start three years earlier. Actually, we have higher notional amounts and they're at higher persistency levels, so that increases the price. If we remove the 200% cap, the price goes up about the same, 20%. That's mainly because your strike prices are increasing. And, finally, what happens when we increase the conservatism in our minimum annuity purchase rates so our ratio runs around 65%? This affects both notional amounts and strike prices. You get a huge cost savings—a 30% decrease in price.

That's Black-Scholes in a simplified framework. How do we do it at AXA? As I said before, it's somewhere between capital markets and an actuarial appraisal method. It's actually in the actuarial appraisal method where we discount our profits and we generate revenues, claims, and expenses, and set up our reserves. We set up our capital, but the new thing we do is explicitly reflect capital markets hedging. We reflect the cost of the option, the payoff of the option, though we never hedge 100% of the risks because we've mapped out this risk a million different ways, and we found that the only way to make money is to either charge a million basis points or hedge a little bit less. But we do it prudently. We never take unsafe risks. Some companies want to offer annuitizations at any point during the year and those are risks that we won't participate in. The next thing we do is have a risk-management plan in place that calls for constant monitoring of the risks. You have to keep track of your book before it moves on you. We mitigate the risks either through hedging or through retrocession and security design, but finally we stress-test all key variables, such as the annuitization rate, just to determine the impact on capital.

The last thing we're going to do is look at some of the key assumptions in our example and how we handle it at AXA.

Persistency

This is right in the notional amount formula. In our Black-Scholes example we assumed a constant 6% lapse with the mortality, so essentially we have a notional risk here. In reality, lapses and mortality vary with duration, age, and market performance. At AXA Re we dynamically calculate persistency to mitigate understating persistency. For example, we may slow down lapses as the expiration of the waiting period approaches in scenarios where there is poor market performance.

Annuitization

This is also in the notional amount formula. In our example we assumed the level 10%, so here is also a notional risk. In reality annuitants need some sort of incentive to give up liquidity and control of their money so they probably won't exercise the option when it's just in the money. We've actually created an interesting function at AXA where it depends upon age and "in-the-moneyness." We do generate pretty high levels of annuitization, though we never assume 100% occurrence. The ratio, the minimum rates to current rates, in our example was fixed, and that means our interest rates and our mortality were set in stone. But as you know, that ratio is both in the notional amount and in your strike price formulas, so this leads to both a notional and a strike-price risk. What we do at AXA is stochastically generate that ratio and set it at a level that we feel comfortable with. With all these variables they are always stress-tested just to determine the impact on capital.

Next, the fund class. It was in S&P only. In reality, what happens is that these funds are invested in a basket. You always have a basis risk, and the underlying funds are not perfectly correlated with the S&P. What we do at AXA is we adjust our strike prices and our notional amounts to reflect this risk. However, when you get right down to it, we feel a lot more comfortable with the basis risk that we're taking than some of these other risks that are on our books.

And finally, yes finally, we finance the option premium. In our example we calculated the 60 basis points assuming some sort of expected return. Some constant rate parameters such as lapses were set at 6% and so on, but what happens in reality with poor market performance is that your revenues decline. Early lapsation causes your revenues to decline so you have a revenue risk and you have to finance your up-front hedge. Insurance companies are actually in a better position to handle this sort of risk. They can actually increase surrender charges to handle early lapsation. We try to do similar things at AXA Re. There's actually a reinsurance premium. This could be based on the income-benefit base or when they try to participate in some of the surrender charge revenue. Less often we may create a revenue hedge where we pass up some of the good performance in order to receive some extra premium on the downside.

CHART 1
MAIN RISKS – ILLUSTRATION

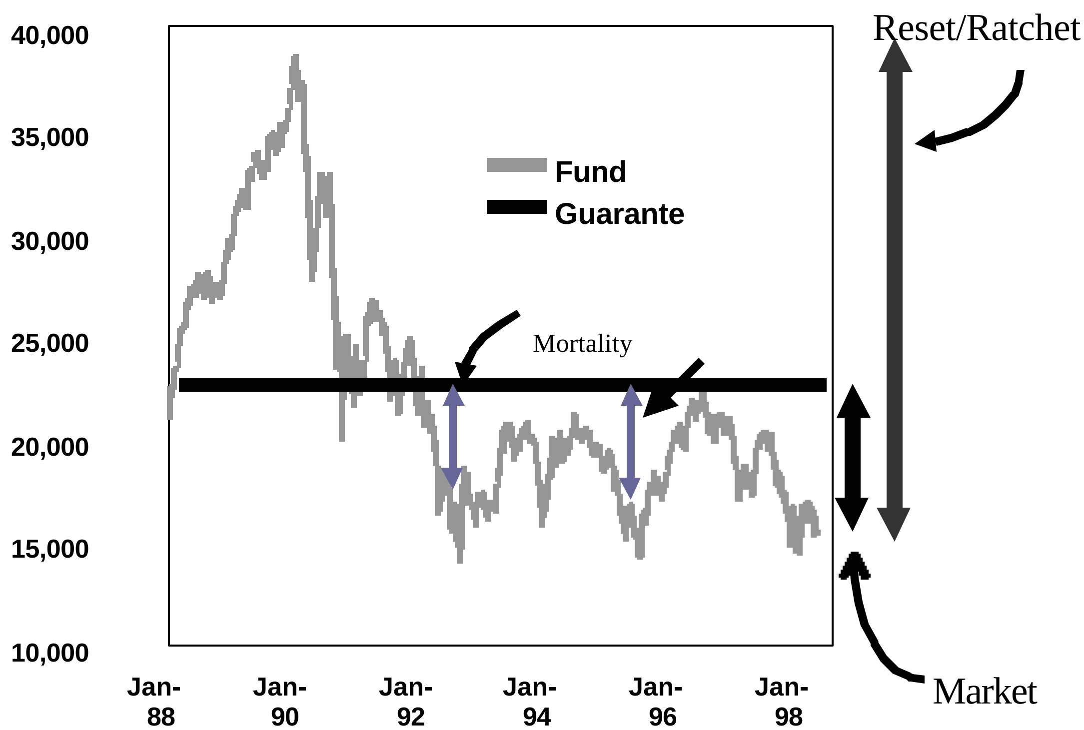

CHART 2
SEGREGATED FUND PRODUCT

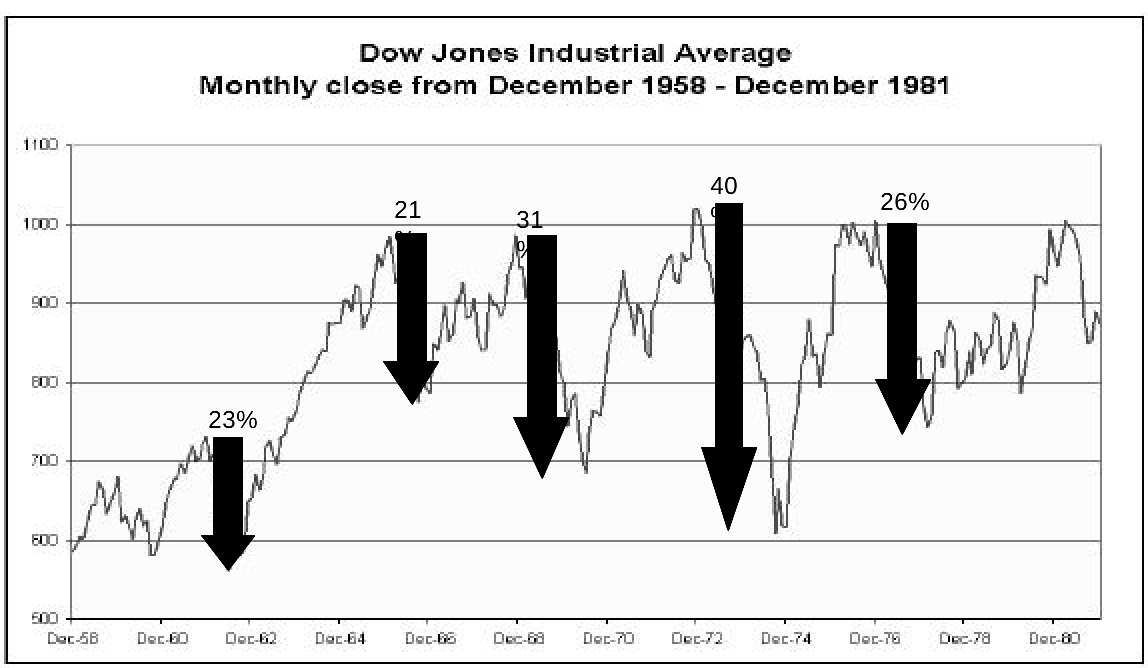

CHART 3
RISK MANAGEMENT—HEDGING COST

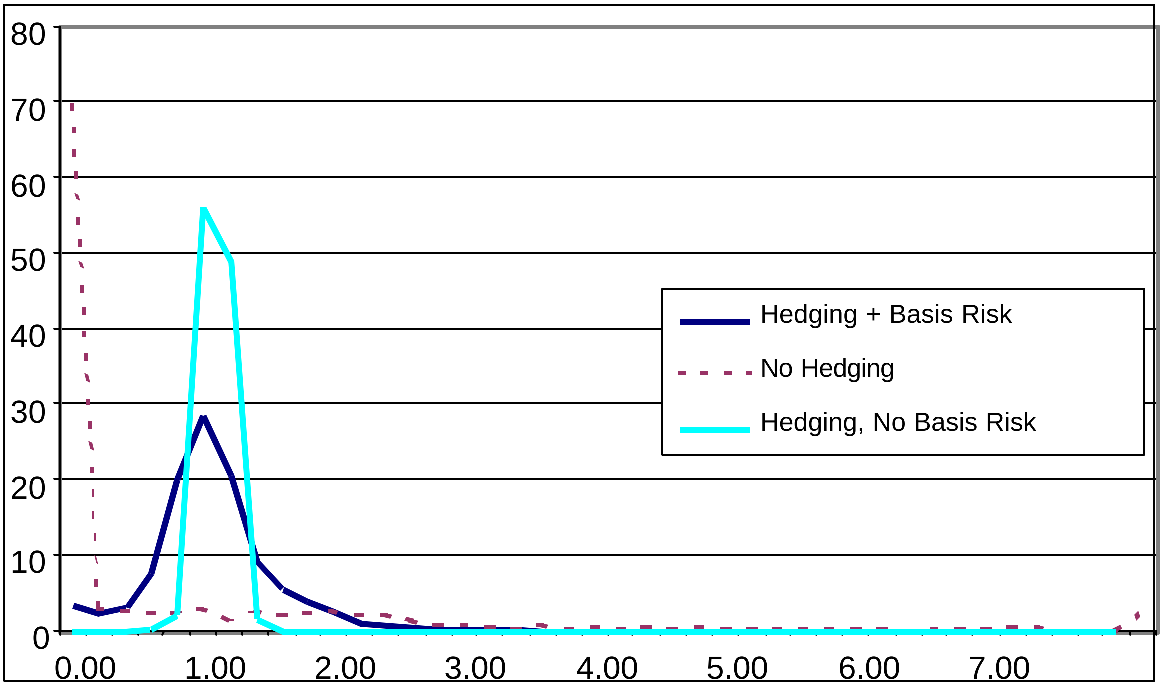

CHART 4
GMIB COST USING BLACK-SCHOLES

CHART 5
GMIB COST USING BLACK-SCHOLES