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

DERIVATIVES IN AN INSURANCE CONTEXT

Moderator: JOSEPH D. KOLTISKO
Panelists: BRYAN E. BOUDREAU
JOHN F. MULHOLLAND
Recorder: JOSEPH D. KOLTISKO

Panelists will discuss practical uses of derivatives, including interest rate and equity options and swaps.

MR. JOSEPH D. KOLTISKO: Derivatives are in the news a lot these days. I'm sure all of you are concerned about how they're going to impact your B1 and C1 risk. Bryan Boudreau of MetLife and Jeff Mulholland of Goldman Sachs will address that issue. I'm from Chase Manhattan. We will talk about some of the derivative contracts that you might find useful in an insurance company product development context.

But first a disclaimer: the things we're going to say are just our own views. They don't reflect necessarily on the policies or practices at any of our firms.

Back to risk, B1 and C1 risk are the risk that your company ends up on page B1 or C1 of *The Wall Street Journal* because of the use of derivatives.

MR. BRYAN EDWARD BOUDREAU: I will talk about insurance company derivative applications. I'll split my comments into two main parts. The first part is an overview of liability risks. This will concentrate not so much on mortality and morbidity risk, but rather on fixed-income risk, which is generally where derivatives are used. Then I'll present two case studies, which are things that we've actually looked at at MetLife. The first is hedging disintermediation risk on a newly developed universal life product. The second is an analysis of some equity participation notes.

There are a wide variety of fixed-income liability risks, but it's convenient to look at most of them as one of two basic types: dollar duration and option risk. Dollar duration risk is two-sided and option risk is one-sided.

Dollar duration risk, as the name implies, arises when the dollar durations of assets and liabilities don't match up. Obviously, that could happen if the actuary tells the investment portfolio manager that the liabilities have a duration of three and the portfolio manager has a view and runs the portfolio at four. In that case, you probably won't do any kind of hedging. But dollar duration risk also shows up when you have fixed forward transactions. For example, suppose you sell an annuity based on current rates, but the premium is not going to come in for a month. Here you have risk because you don't know what rates will be when you invest the proceeds. Another example of this type of risk is in the LTD line. Premiums are set by using discount rates that assume a certain fixed investment return, but the premiums are effective for a period of several years. Again, there is a risk that rates may drop significantly before premiums could be adjusted.

You can measure dollar duration risk in two ways. You can look at yield curve level risk and measure total duration mismatch. You can also look at yield curve shape risk. Here you'd employ a partial duration approach such as the one Bob Reitano frequently talks about. In either case, however, the nature of dollar duration risk is that there are

approximately equal exposures to gain or loss. You'll gain if rates move in one direction and lose if they move in the other.

What kinds of strategies might be used to hedge dollar duration risk? In general, you want to add or subtract dollar duration to get your portfolio back to a neutral position. One way to do this is via forward purchases of assets. If you've written a liability forward, then buying an appropriate asset forward adds the dollar duration required to get your portfolio matched.

Futures transactions and fixed-to-floating swaps can also be used to add or subtract dollar duration. If you're concerned about different points on the yield curve, you could also use floating-to-floating swaps. For example, you could receive floating coupons based on the one-year rate and pay floating coupons based on the five-year rate. All these hedging strategies have payoffs that are two-sided. You may win or lose on the hedge, depending on interest rate movements, but you hope you're flattening out your performance profile.

In contrast with dollar duration risk, insurers may also be exposed to one-sided option-type risks. Insurers may be short options on either the asset or the liability side of the balance sheet. A familiar example of being short an option on the asset side would be when you've purchased mortgage-backed securities and you've implicitly written call options.

The insurer may also be short options on the liability side. An example of this would be book-value disintermediation risk. If products offer book-value surrenders but the assets can only be liquidated at market, there may be some exposure to rising rates. Rate windows and minimum guarantees also represent options. If the insurer holds a rate open for a period of time, you run the risk of receiving a large deposit when rates fall. But this is one-sided risk exposure. If you're short an option and rates move in one direction, it's not going to affect you. But if rates move in the other direction, then you're going to lose.

So what kinds of strategies might be used to hedge convexity risk? Well, if you're short options, then you might want to buy options back. In general, you'll pay some sort of premium in exchange for a potential payoff if rates move in the direction where you have exposure. Examples of these types of hedging strategies would include options on futures contracts, which pay off only if rates move in one direction. Similarly, interest rate caps and floors pay off only if rates move above or below a certain strike level. And there are also interest rate swaptions, which give you the right to enter an interest rate swap contract. In each case, the hedge instrument pays off only if interest rates move in a certain direction.

The first case study I'd like to look at is hedging book-value disintermediation risk. This is an actual case that we've modeled fairly extensively. The product is corporate-owned universal life with a sizable initial premium. The credited rate is reset annually. There is a minimum guaranteed rate. There are also book-value surrender rights.

These product features combine to give financial exposure to rising interest rates. If rates increase, then one of two things could happen. You may need to increase your credited rate faster than your asset book yield is going up in order to keep the money.

Alternatively, if you don't move the rate up quickly enough, you may get lapses, in which case you're going to take market-value losses on your assets. Your exposure is basically one-sided.

In order to model and analyze the risk, you must pin down certain assumptions. First is the crediting rate strategy. You need to model what you intend to do with rates under alternative scenarios. The lapse assumption is probably the most crucial determinant of how much risk you have. You need to test a base assumption and also test for sensitivity to different lapse rates. You also need to model the investment strategy. The objective is to see how the credited rate strategy plays out when combined with the investment strategy.

We ran this product through an asset/liability model (Chart 1). We varied the level of interest rates and measured the return by using the present value of statutory profit. The base rates were as of last year, when the five-year rate was around 7%. Without hedging, the performance profile is consistent with what you'd expect. What you see as rates move up significantly is a decrease in the present value of profits. The drop in the performance profile for high rates is the effect of disintermediation. The bend at the other end of the graph, when rates fall, is the effect of minimum guaranteed rates.

CHART 1
PERFORMANCE PROFILE WITHOUT HEDGING

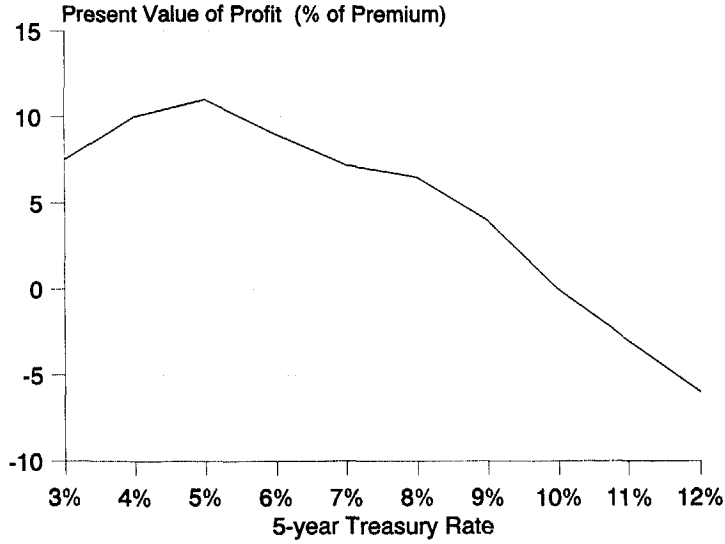

We looked at constant maturity Treasury (CMT) caps as a hedge for the disintermediation risk in this product. In exchange for a determined premium, CMT caps pay coupons equal to a notional amount times the amount by which a designated index rate exceeds a specified strike rate. If the index rate is less than the strike rate, the caps pay nothing.

We looked at cap hedging strategies with a variety of terms, strike rates, and notional amounts. What we came up with as a reasonable hedge strategy was a CMT cap with a notional amount equal to 80% of the premium and a five-year term. The cap payments were determined annually by using the five-year CMT rate, with a strike rate 150 basis points over the spot five-year CMT rate.

The effects of hedging can be seen in Chart 2. The sloped line is the same unhedged performance profile from the previous graph. But now we've also included the profile for the hedged portfolio. You can see that if interest rates stay the same or move lower, you'll sacrifice some profitability. But if rates move higher, the performance profile is now more favorable for the hedged portfolio. So what you've done, in effect, is purchase insurance, paying a premium in exchange for protection against rising rates.

CHART 2
PERFORMANCE PROFILE WITH HEDGE OVERLAY

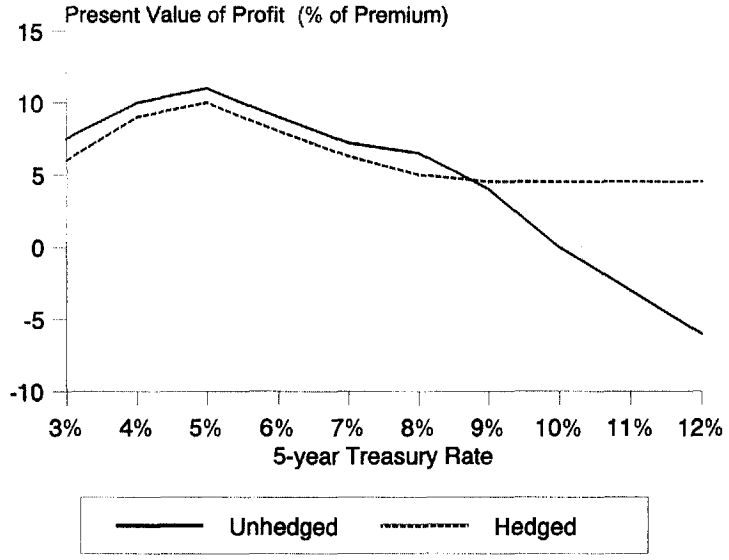

Modeling assumptions are important. You must test the sensitivity of your modeling results to different assumptions. In particular, you'll want to look at different lapse functions. In addition, you'll want to look at different types of interest rate scenarios.

In addition to modeling, other issues should also be addressed when considering derivative hedge strategies. One issue is the notion of risk tolerance versus option cost. Is the premium reasonable in relation to the risk reduction? The answer to this question is a function of risk tolerance. You need to know the effects of hedging on profitability and determine how much profit you're willing to give up to reduce risk.

You'll want to consider the nature of the risk. Is the risk one-sided or an option-type risk? You want to consider the effects of diversified product portfolios on the aggregate

level of company risk. An insurer may have one product line with offsetting risks that acts as a natural hedge to another product. You need to keep in mind the determinants of the hedging cost. Derivative instruments are priced based on the forward curve and interest rate volatility. Because these parameters are changing constantly, the cost of hedging with derivatives will vary over time. And of course you'll need to keep in mind the regulatory issues.

The other case study I want to look at is equity-linked notes (ELNs). These instruments have been designed as alternatives to direct equity holdings, given the relatively high risk-based capital charge on equities. These notes specify a minimum return plus equity upside participation. For example, you might receive the greater of 5% compounded to maturity, or x% of the return of the S&P 500 index, where the x% varies depending on the term of the note and the yield curve at the time of pricing. The notes are considered bonds for risk-based capital calculations.

Equity-linked notes are basically hybrid instruments. In effect, the note can be analyzed as a combination of a zero-coupon bond plus call options on an equity index. Alternatively, you could view the note as an equity holding, plus a put option that guarantees a minimum realized return.

We recently analyzed these notes. What we were looking at was a "menu" of S&P 500 linked notes. We wanted to determine if any of the structures made sense as alternatives to holding equity directly. ELNs are available from different issuers in a variety of maturities, ranging from 3 years out to 10 and even 20 years. You can get different minimum returns and equity participation factors. The type of equity participation also varies. The equity payoff may be a function of either the price appreciation or the total return of the S&P 500 index.

I want to describe some of the concerns that are important when analyzing these types of structures. It's important to focus on the economics. The note should be analyzed as the sum of the parts—in this case, a call option plus a zero coupon bond. The options in these notes can often be analyzed by using straightforward Black-Scholes techniques. However, for complicated payoff structures, you may need to use random scenarios to price these options.

You'll also want to look at total return distributions. You'll want to analyze the performance relative to bonds and direct equity holding. Because the returns for these structures typically aren't symmetrical, you may also want to look at downside risk measures, that is, the distribution of returns below a certain level. The effect of taxes is also important.

When we performed our analysis, we found that using risk-based capital (RBC) adjusted returns can be misleading when comparing equity linked notes with stocks and bonds. So it's important to compare economic results on a consistent basis. And not surprisingly, we also found that the structures had total return performance characteristics in between stocks and bonds. Results are sensitive to assumptions, particularly the distribution of equity returns, so you should test for different return distributions. And lastly, the tax effects are significant. You get significant tax deferral under current tax rules, and that tends to favor longer note structures.

MR. KOLTISKO: I want to amplify and give some background material on the same topics. Often we look at derivatives as if they're some sort of magical, powerful weapon, something such as fire or nuclear energy. That's the wrong way to look at them. Rather, they're just another set of pliers in the toolbox to use in managing an insurance company's investment risk.

I think it's not an effective or useful frame of mind for someone managing the investment portfolio of an insurance company to make a black and white distinction between risky and risk-free. Insurance companies are already in the risk-taking or risk management business. It's impossible to avoid risk completely. Derivatives allow portfolio managers to sell off unacceptable risks and (in a controlled way) add certain equity or commodity exposures to otherwise plain vanilla general account portfolios. That's particularly true, for example, in the group annuity area, and in the corporate-owned life insurance (COLI) area, and in other niche markets where there may be an attraction to a sophisticated buyer for upside potential.

Actuaries have a sort of morbid fascination with risk, death, catastrophe, asset default, and interest rate volatility. But not one of these risks has a right to dominate the performance profile. I'm sure you're aware of the forms that investment risk takes, including liquidity risk, prepayment and extension risk for collateralized mortgage obligations (CMOs), as well as equity and foreign exchange risk.

How do you enforce the view that you're not going to let any one particular risk factor take control? Let me answer this by way of example. Suppose your sales representative calls you with a great product such as a bond that yields 600 basis points over Treasuries. As a sophisticated investor, you're going to ask what the yield is going to be tomorrow and will you get your principal back and when and how. Clearly, the distribution of returns is more important than today's yield. Now broaden that to two or more securities. If total return losses on one security will exactly offset total return gains on the other security, basic portfolio theory tells you that risk of the portfolio is reduced. These concepts of risk and return are familiar.

How many of you are familiar with the term value-at-risk? This is a portfolio management method that has grown into a program used to manage traders. In this technique, we determine a probability distribution for the value of the portfolio based on recent historical returns in certain asset classes and on the correlation among those classes. In a trading context, this is used to manage traders by determining a confidence interval for the amount of possible losses the firm may incur from day to day. Management first decides on an acceptable chance, say a 5% chance of a loss above a given dollar threshold. The crux of the method is to monitor and control that threshold amount, which varies from day to day. Whether 5% or 1% is the appropriate and useful benchmark is a judgment call. The observed probability distribution that I described earlier is affected by the trading activity. If the dollar amount of the threshold gets too high in relation to the firm's capital, it will institute a hedging program (or fire the trader). The goal of our corporate risk management analysis should be to come up with a distribution function that helps us analyze our risk a little better to decide where our losses are and sell them off. That's what we'll see in the case study.

First I have some vocabulary and background on over-the-counter (OTC) interest rate contracts. When buying an OTC cap or floor or entering into a swap, you're working

with a counterparty. That's who assumes the other side of the contract. You're paying a premium or entering into a contract with someone to make contingent payments, so you're concerned about your counterparty's credit. Currently, the annual statement reporting for Schedule DB explicitly tries to capture that risk with its potential exposure measure. The contract is going to take place over certain time periods, the tenor of the contract. During that period, contingent payments are going to be made and the payments and the cost of this contract are going to be made with respect to a notional amount. If you talk about a $10 billion swap, it's not as if $10 billion is trading hands. That amount is never at risk. That's simply a reference amount that is used to determine how much the contract costs and the size of the payments to be made from one side to the other of the contract. Those payments are made with reference to a certain strike rate.

Here's an example that Bryan described aptly. In an interest rate cap, initial premiums are paid by the insurer to the swap dealer, so the counterparty here is the swap dealer. Contingent payments come back, represented by the dotted line in Chart 3, based on where, in this example, the five-year Treasury is. If the five-year Treasury is greater than the strike rate of nine, that difference is going to be paid back to the insurer. I think of a cap as an insurance policy for interest rates. If interest rates do hit 10%, you'd receive the difference, or the 10% minus 9% for five-year Treasury, multiplied by that notional payment.

CHART 3
INTEREST RATE CAP

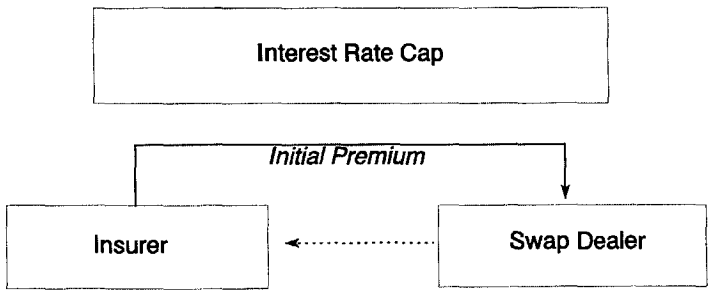

Payment:
5-year U.S. Treasury—9%, if greater than zero;

Net rate multiplied by notional principal

Chart 4 shows a swap. The fixed rate is set at issue and the floating rate is based on current levels. Initially, this should be a zero net present value exchange so that the insurer is paying (in this example) 7.75% on the notional principal. The swap dealer pays back the then-current five-year Treasury. The five-year Treasury today is about 6.75%. Initially, the insurer is making a payment to the swap dealer, but based on where the forward rates are, you'd expect that in the future the swap dealer would be making payments back to the insurer. No cash changes hands initially.

CHART 4
INTEREST RATE SWAP

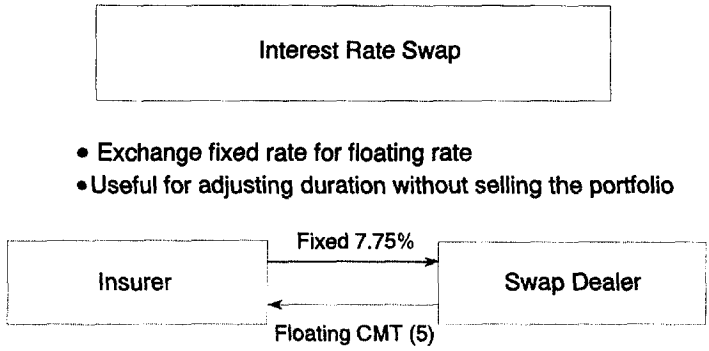

Payment:

Stated rate multiplied by notional principal

What happens if rates rise sharply and unexpectedly above the forward rates? The swap dealer makes higher payments back to the insurer and the insurer ends up gaining. The contract would increase sharply in value. If rates fall precipitously, the contract decreases in value. The insurer might end up paying much more back to the swap dealer than it receives from the dealer. Now what's so scary about this? How can you be in the insurance business and not take risks such as this? The cash flows depend on the level and path of rates. This is the same sort of business that we're in, and that's why I think that it has a place in a controlled way, in a managed way, in our portfolio.

Now for adding risks back to your portfolio, one can invest in a structured deposit, which is just a certificate of deposit with an unusual interest rate. With a CD, of course, principal is returned at maturity. If it's a small amount, FDIC insurance may apply. What's novel about it is that the coupon can be linked to an equity index, a tradable commodity index, or an equity index that gives you an exposure that might be expensive to introduce otherwise.

Equity-linked notes are typically public debt produced through a medium-term note program. The issuer pulls a note off the shelf and issues debt with a coupon that's attractive to you, the buyer. Its coupon is stated in terms of the movement in an equity index or a commodity index or a foreign index. The issuer will turn around and swap it into a rate that is below its cost of funds, so everybody wins with such an approach. The coupon on such notes can be linked to a basket of indexes, rather than just to one index. It might be linked to several indexes that are important to you. The big advantage to the investor really is consolidation. Rather than trying to hire a manager to invest in commodities and incur the related management, research, transaction, and compensation costs, instead you'll have a note that achieves a similar payout. The accounting and the costs are consolidated in one statement. There are many hidden costs to having an equity manager or a portfolio manager in a small company. You might find it more advantageous to get that exposure through an equity-linked principal-protected structured note.

My understanding is that as long as principal is returned at the end of the period, those are treated as bonds as far as the Model Investment Law is concerned, and they receive the bond RBC factors. The other big advantage is speed. If you want to implement an asset allocation strategy, equity-linked notes allow you to get that exposure to, say Belgian equities, quickly and cheaply.

We talked about some of the ways to sell off certain risks and add certain risks. To use derivatives rationally, we need a probability distribution for firm value. We have some tools that can get us there, and we use those tools in cash-flow testing. But, we've all been involved in wheel spinning when it comes to cash-flow testing, of projecting flows that aren't related to what an investment manager will actually do. I would argue that effective cash-flow testing has to relate the firm's economic value to the market value of tradable assets. If it's on an inconsistent basis, there will be no way to add or subtract certain assets to modify the risk profile of the company. Then I don't see how you implement all those insights you gain from cash-flow testing.

There are other practical constraints, of course, that come up in trying to run a reasonable, theoretically correct, dynamic solvency test. There is the choice of interest rate scenarios. You may need 10,000 or 20,000 possible paths of interest rates to come up with a reasonable risk profile of the security and equally so for the firm. If you have a detailed insurance model, though, five or ten scenarios may take all night to run on your system, and that makes it impractical.

Rather than a stochastic approach, what we've used in this study is a probabilistic approach, which picks certain interest rate scenarios and assigns a probability to that scenario. Once you have a scenario, it's a simple matter to project the cash flow of certain interest rate swaps or caps or floors, discounting them today and coming up with a market value for the contract under that particular scenario. We've done the same sort of thing for the company, and this example that we'll talk about is based on a particular firm.

MR. ROBERT A. NELSON: How do you assign probabilities to scenarios?

MR. KOLTISKO: What we've done in this example, as a practical step, is to reverse the order of taking expected values. Instead of getting the expected value of the firm as the average discounted cash flow after running many, many scenarios, we choose scenarios based on historical ranges about the forward curve at a given horizon. I think that's an acceptable approach, and it's often used as an approximation to manage portfolio trading risk. The baseline is the forward path of interest rates. Why in cash-flow testing do we use a level scenario? How likely is that? By contrast, the forward rates project the market's expectations of where interest rates are going to go and reflect the cost to hedge. That's an effective basis for deciding where the value of the firm should be.

In this example, what's been done is to take the path of forward rates, look at a reasonable horizon (here one year), apply an assumed distribution (here lognormal) for where that rate will be in a year, and finally draw a scenario that gets you to where rates should be one year out. Both CMT 10 and the London Interbank Offered Rate (LIBOR) were projected separately. Of course, a key input is the expected level of volatility. I look one and two standard deviations above or below the forwards to form the basis for the scenarios that are used in this analysis.

The assumed cumulative distribution function for the rate gives us the likelihood for each scenario. We assign the likelihood of that scenario to the likelihood of the present value of earnings under that scenario. The answer to your question depends on which interest rate distribution function is used and whether historical or best-estimate parameters are applied.

The value of the firm is measured according to the economic value school of thought. Rather than looking at statutory earnings or simply at GAAP equity, you roll forward the statutory entity, keep whatever target surplus is required to stay in business as an insurance company, and dividend out to owners what will be available. Where it differs from a traditional economic value, though, is in discounting those dividends at the short-term rate under the scenario. The goal is to get a discounted cash-flow value that is in some sense consistent with current market values. If you're an equity investor looking to buy the company, you might discount these as 12% or 15% or whatever your hurdle rate would be. That 500-basis-point differential is your risk premium for buying the company. But as a manager, you're faced with short-fall risk. You're concerned with losses that may occur and you're going to fund those shortfalls by borrowing. We therefore discount at the short-term rate at which you would have to borrow to fund losses. Based on that, if we can project a range of scenarios and come up with market values of the earnings that are dividended out, we end up with a distribution for the value of the firm as in Chart 5. There are ways to improve the methodology and the illustration, but unless we can somehow capture this, I don't think we're being very effective in cash-flow testing.

Chart 5 should probably more appropriately be a bar graph. Instead, I've connected the dots from the ten scenarios used to illustrate it a little more clearly. As you can see, the expected value for this distribution is about $2 billion, and that should relate reasonably well to the equity of the company. But there's a significant chance that firm value is going to be less than a billion dollars or even less than zero. As an asset/liability modeling (ALM) manager, you should decide how to hedge that risk to maintain and maximize the value of the firm.

The next few charts show how overlaying certain caps and floors affect that distribution. Chart 6 shows the effect of caps. In the center are the forward rates. Moving up and down we are movements up half of one or one-and-a-half and two standard deviations away from where the forward rates are going to be. The solid black line in Chart 6 shows what happens to the firm in those movements of interest rates. For a plus-two standard deviation movement, it becomes insolvent. But when rates fall, as we know, negative convexities and all the other options in the insurance products tend to limit the upside potential. It's not like a zero coupon bond. As rates fall, we lose some of that upside. The dotted lines in Chart 6 show the effect of caps. We used a straight 9% cap and a forward-starting 9% cap.

In the low rate or flat-rate scenarios, we've paid an option premium and nothing's come back to us. Rates have fallen and we end up with a lower value. On the other hand, as rates rise the cap pays off and protects us in that very unlikely +2 standard deviation movement up. The break-even point is at about a one standard deviation movement up in rates. At that point you're just as well off whether you bought the cap or not.

CHART 5
UNHEDGED FIRM VALUE
PDF

Mean : $ 1.98 Bn.
Std Dev: 775 Mn.

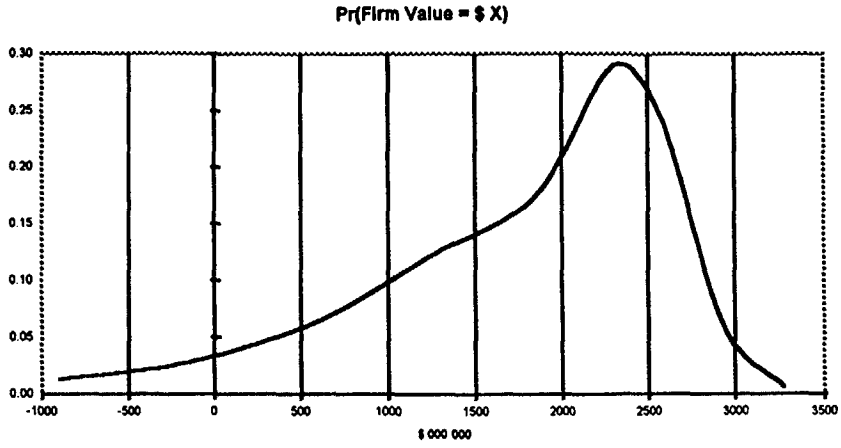

CHART 6
CMT(5) CAPS

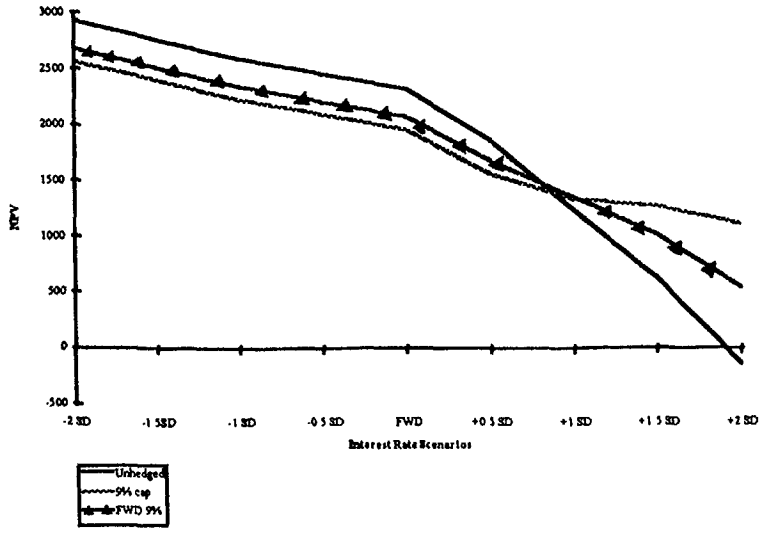

Chart 7 expresses the same point in a distribution graph. We have the same distribution for firm value that was there before, then overlaid on that are the probability distribution functions (PDFs) for the hedged firm value with the cap. What happens is that we've cut off the losing left side. That should be an attractive contract for someone looking at controlling the risk in the firm. That comes at a cost, though, because now the mean value is lowered. You can see that distribution has a peak at a lower level, and that's because you've paid an option premium for a contract that doesn't pay off in the future.

CHART 7
CMT(5) CAPS
PDF

Notional amt: $15.4 Bn.

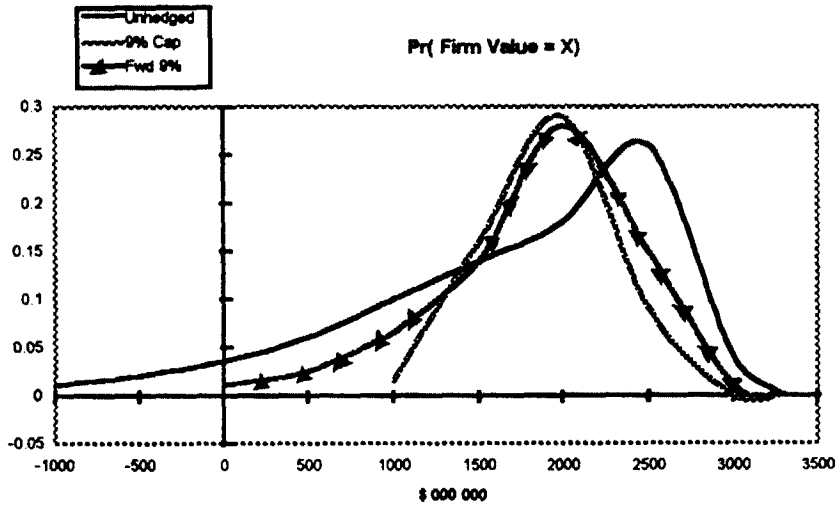

MR. NELSON: As I move along the PDF, I change the economic value. Try to think of this in terms of the market value of a firm measured by the stock price, and just that stock price. It would seem to me that as I hedge it enough to get it to look just like a ten-year bond, my stock would trade as a price/earnings (P/E) ratio, such as a ten-year bond. Bonds pay yield, but in the end, people would just be buying my company for a dividend, which is the same thing as the coupon. I could end up by having my stock priced as a 10-year bond, which is not what equity investors want. You want to hedge some risk, but you don't want to get to the point where you're trading like a ten-year bond. Have you thought about how to maximize the stock value? What kind of shape looks best?

MR. JOHN F. MULHOLLAND: We'll come back to that. For what it's worth, I'll try to tie in their two talks and try and put it within the framework where you're headed.

MR. KOLTISKO: In Chart 8, we've put in a swap and, again, the heavy line is the unhedged value. The lighter of the two light lines is the cap from before. Again at the forward rates, you lose your premium by buying the cap. By contrast, the swap gives a

lot of the same protection. At the forward rates, the value of the swap is neutral. It's a zero net present value exchange at issue, so the value of the firm is identical. As rates rise, the swap makes a payment and it improves the operating performance. As risks fall, however, you're paying for that protection by significantly reducing the earnings that can be achieved.

CHART 8
CMT(5) SWAP

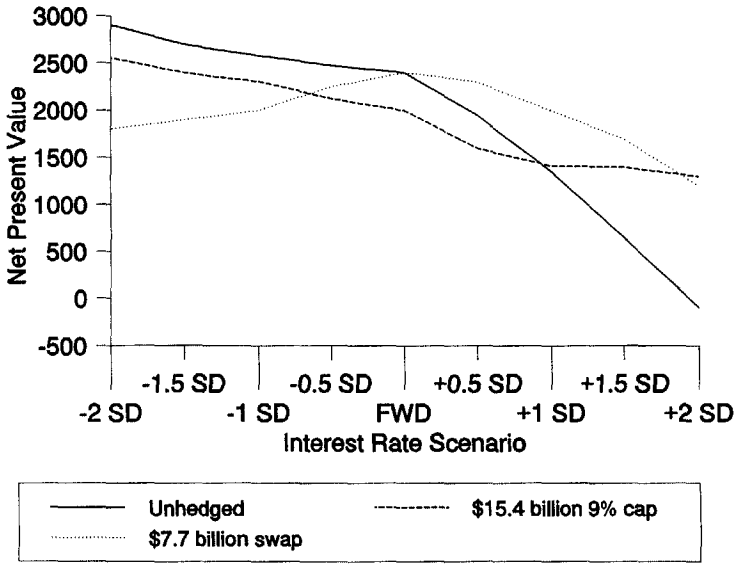

One way to limit that is to buy a floor contract with the swap, which is what's called a participation. For a smaller premium, that allows you to level or take away some of that unacceptable downside. The effect of this is to rotate the risk profile in a way that can be attractive. Its attractiveness is shown in Chart 9, which just shows that rather than having a distribution that's wider, what we've done is focus it more tightly so the standard deviation of returns is significantly narrowed. The downside has been sold off so you're more confident that the firm is still going to be worth its $2 billion expected value.

CHART 9
CMT(5) SWAP
PDF

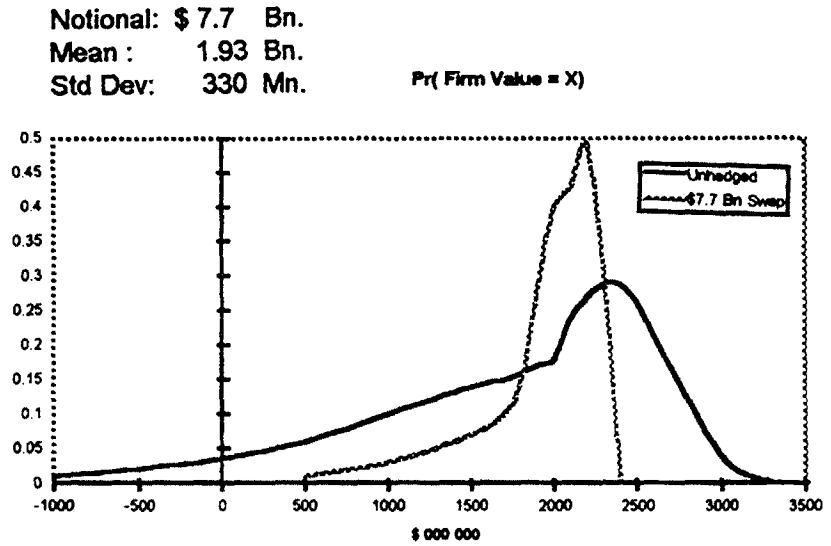

Other issues relate to this analysis. There is the issue of whether the shareholders pay for this derivative contract or whether the policyholders do. If it's bought as an asset in the portfolio, the policyholders are going to be paying for whatever option costs are there in an unfavorable scenario through their crediting rates or through the investment returns. If it's purchased by the shareholders, perhaps as a separate entity, perhaps in a funding subsidiary, those interactions don't take place. These interactions are worth analyzing in some detail, but it doesn't obscure the fact that this approach seems to improve the profile.

In conclusion, this sort of decision process helps reduce the standard deviation, the risk, and helps analyze the mean for the value of the firm. I want to leave with you the thought that you can manage these sorts of risks by using caps and floors and other interest rate options. Further, there are significant advantages, regulatory and simple risk management advantages, for getting those desired exposures through equity-linked notes or structured deposits and other packaged equity-like products.

MR. MULHOLLAND: I actually have prepared slides; however, I will set them aside because you have raised an extremely interesting question, and I would like to address the issues discussed by the previous two speakers within the context of the irrational world in which we live.

Let's examine a case study and discuss some of the irrational characteristics of the insurance company business environment. However, first we must make several assumptions:

- We are employed by a stock company.
- Management's bottom line is GAAP income.

- The company is a monoline single-premium deferred annuity (SPDA) writer.
- The asset /liability management system takes approximately a week to run 20 interest rate scenarios.
- The interest rate scenarios that are run to measure the business are centered around today's yield levels (a level interest rate scenario).

Many companies' actuaries may find this risk management situation familiar; however, we would make the case that some important biases will result from the analysis if conducted from a traditional actuarial perspective. We would like to recognize, though, that this is a very difficult problem that may be addressed from many different perspectives and entail working within each company's computer resource constraints.

Our basic premise is that we believe that assets and liabilities should be valued consistently from an interest rate perspective. Also, the level at which a company buys and sells assets on a daily basis should determine what the appropriate expected values (forward rates) and standard deviations (volatilities) of the future distribution of interest rates should be. (We can deduce what investors' best estimates of these parameters are in the aggregate through the relationships among different securities and options traded in the fixed-income market.) We do not believe that these parameters will necessarily be accurate in predicting future interest rate levels, but they do assist in analyzing the opportunity cost of retaining interest rate mismatches (the up-front cost of hedging).

Traditional actuarial analysis suggests using level interest rates as the base-case scenario. However, the bond market allows investors to lock in higher interest rates in the future without incurring any up-front cost. For example, one may lock in a long (or short) position in five-year zero coupon Treasury rates five years from now by going long (short) a ten-year zero coupon Treasury security and short (long) a five-year zero coupon Treasury security. We do not maintain that implied forward rates are an accurate predictor of future interest rates; however, we do maintain that these are the proper basis for measuring interest rate risk, because they are the basis for measuring value in the instruments that may be used today to adjust a company's interest rate risk profile.

The same analogy can be made for choosing to measure interest rate risk by using market-implied volatilities (the market's estimate of the standard deviation of the future distribution of interest rates). These market standard deviations can be used as a benchmark for determining how likely it is that interest rates will reach a certain level over a given time period (how high is too high). By using this measure, for example, companies can deduce how reasonable it is for interest rates to rise 500 basis points (5%) over the next five years based upon what the hedges cost (and what that cost is telling us about the standard deviation of the future distribution of interest rates).

If, however, a company ignores the fixed-income market's levels in determining its risk, then certain biases result. For example, using a base case biased toward lower rates (because level rates will be lower than forward rates) will produce a longer effective duration of liabilities for our monoline SPDA writer. Many traditional actuarial analyses suggest an effective duration of four for this liability, but others who use market information in their modeling believe it's closer to two (e.g., assumes a seasoned block of business, a career agency distribution force, etc.).

Using scenarios biased toward lower rates also suggests maintaining an asset portfolio with a longer duration than the liability. This conclusion develops because assets that have a higher option-adjusted duration will outperform relative to shorter liabilities, if interest rates fall.

Based on the traditional methodology, many SPDA companies had longer asset than liability durations and benefited from the bond market rally during 1991, 1992, and 1993. However, these same companies suffered greatly during the rise in interest rates during 1994. For example, a significant number of companies failed the pop-up scenario of the New York Regulation 126 at the end of 1994, which assumed that there would be a second year of interest rate increases following the year's bearish bond market performance.

However, because stock companies are typically judged by the equity market based upon earnings' performance, for many companies the rational bottom line may be earnings rather than market value or economics (because the equity market has access to very little information regarding the true economics of the SPDA business). As a result, many companies with which we work have been examining their risk profile by projecting earnings by calendar year and then examining if the results are within an acceptable range. To the extent they are not, then we assist them in examining the trade-offs among various risk mitigation techniques. In all our analysis, though, interest rate sensitivities are analyzed consistently among assets, liabilities, and hedge instruments.

For a typical SPDA portfolio, the worst earnings scenarios usually arise from rising rate scenarios. To hedge this risk, the most common solution is to purchase caps on Treasury rates (CMT caps) or put swaptions (options on swap rates). Option-based hedge programs are more common today (in light of Orange County, Procter & Gamble, etc.), because the worst-case hedge performance is the loss of option premium, whereas swaps and futures expose companies to hedge losses if interest rates fall. Options programs have become preferable because the loss of option premium (if the options do not mature in the money) can be compared with owning fire insurance and not having your house burn down.

CMT caps have been the traditional options-based solution for this hedging application; however, recently the liquidity (or width of the bid/offer) of this options structure has come into question. OTC puts on Treasury notes and cash-settled put swaptions provide significantly better liquidity and structuring flexibility.

MR. LINGDE HONG: I have some experience in ALM work with the New York seven scenarios. Risk was obvious. Intuitively (and recent experience shows this) you can tell that hot money will lapse and there will be losses in a rising rate environment. Why do companies find that surprising?

MR. MULHOLLAND: Many people found out by doing their pop-up tests at year-end 1994. For many of these people it was an epiphany. As I mentioned, some companies flunk the pop-up scenario by several hundred million dollars. Apparently, companies in that situation did a combination of things: they implemented hedge programs or transferred in surplus, depending upon the company's situation. Yes, the New York seven has many problems with it. A couple of the goods things about it are the pop-up and pop-down scenarios.

MR. HONG: To follow up on the pop-down scenario, reinvestment is a longer-term risk. We assume that management will deal with it in a rational way. How does it do that?

MR. MULHOLLAND: Oftentimes, those risks are reviewed as offsetting. Risks with rising rates on one product are assumed to offset risks with falling rates on the other. What's interesting is you're really exposed to rising and falling rates on the same products at the same time. You're exposed to rising rates less on the longer term. If you pull out the duration offsetting and just look at the optionality of the products, to the extent one of the disaster scenarios is for rates to go down, you have reinvestment problems on the long-term product and minimum rate guarantee problems on the annuities. These are very complex issues. In some lines of business you have some problems from falling rates and some from rising rates, and in some products you have exposure to both. You're short a straddle.

MR. HONG: Have people done a lot of work in terms of offsetting hedging for purely long product companies?

MR. MULHOLLAND: Yes, but it varies drastically. Companies are practically defined by their risk tolerance. One company may spend $15 million in option premium every year, protecting itself against reinvestment risk, while another in the same market does nothing.

MR. KOLTISKO: I have a question for Bryan about implementation. We talked about trying to implement these issues on a company basis. How realistic do you think it is to ask the manager of one line of business to enter into contracts, such as interest rate swaps, with other portfolio managers within the company? Do you think that is a practical way to allow for the correlation and covariance that you mentioned within the product line?

MR. BOUDREAU: Yes, I actually think it has a lot of hope. It's something that comes up often enough at our company. I think it would be an effective way, in fact, to transfer risks between lines of business. If you have got a lot of business that has exposure when rates go up and you have another line that has exposure when rates go down, why not write a swap, essentially, between the two lines of business, transfer the risk internally, and save the amount that you'll be paying to Wall Street to buy both sides? I think that companies probably will start to look at that sort of thing. I think there's a lot of hope for it.