RECORD OF SOCIETY OF ACTUARIES
1989 VOL. 15 NO. 2

THE INTEGRATION OF INVESTMENT FEATURES INTO ACTUARIAL ANALYSIS

Moderator: MARK WILLIAM GRIFFIN
Panelists: D. ANDREW HALL III
ROBERT L. WHALEN*
Recorder: MARK WILLIAM GRIFFIN

- o The actuary should be aware of techniques that are available to compare the expected returns from many different types of investments. Understanding these techniques will assist the actuary in pricing many of today's insurance products as well as measuring their ongoing profitability.
  - -- Call protection versus rate of return
  - -- Default risk, marketability, and maturity structure versus rate of return
  - -- Pricing example
  - -- Different products
  - -- Application to valuation

MR. MARK WILLIAM GRIFFIN: Our first speaker, Bob Whalen, is Vice-President and Director of Investment Technology at CIGNA. Our second speaker, Andrew Hall, is a Vice-President in Morgan Stanley's Mortgage Backed Area. I'm Mark Griffin, the third speaker, and I'm part of Morgan Stanley's Insurance Coverage Group. Before Bob gets started, I'd like to make one introductory comment. In a nutshell, we hope to give you a closer look at the asset side, and also make you aware of asset side areas where actuarial analysis is, or could be, done. At the same time I hope to show where an actuary might adopt some asset side techniques into his or her work.

MR. ROBERT L. WHALEN: In my 30-plus years at CIGNA and its predecessor company, Connecticut General, I had the pleasure, at least most of the time, of dealing with many actuaries both inside and outside the company.

Over that time, I've developed a considerable admiration for the skills and techniques that actuaries bring to their tasks and, especially in recent years, an appreciation of the difficulty that the current markets, both investment markets and insurance markets, are making or bringing to the tasks the actuaries have to do. I'm going to be talking about how cash flow arrays, and that in plain English means assets and liabilities, behave in environments of changing interest rates. Now, lots of the concepts that I'll be talking about will be familiar to all of us. I'll mention some of these in the interest of having an orderly and complete discussion.

To begin with, what cash flow arrays are we concerned about? On the asset side, I will talk about fixed rate investments, bonds, commercial mortgages, pass-through and the many variations that have come along in the last few years, collateralized mortgage obligations (CMOs), interest only pieces (IOs), principal only pieces (POs), etc.; in fact, any fixed rate, fixed maturity investment instrument.

I'm not going to be talking about equity instruments. Equity instruments' behavior is affected by too many factors for their behavior to be predicted with anything like the certainty of bonds and so forth. We'll also be talking about actuarial liabilities, and in particular the liabilities connected with group annuities, with GICs, and with any actuarial liability set whose time and dollar incidence is certain; certain meaning in an actuarial sense. For example, take the expected value of mortality. I'm going to be laying on you a number of principles of my own devising. Most are not particularly original, certainly not original with me. I'm sure that all of you are delighted that you've come 3,000 miles to find out principles number one and two: 1) when rates fall, values rise; 2) vice-versa number 1. I'm constantly amazed, quite honestly enough, among the laymen

* Mr. Whalen, not a member of the Society, is Vice President of Investment Technology at CIGNA Corporation in Hartford, Connecticut.

amongst us, that people really don't understand this, particularly in periods like just now. They are dismayed to find that their bond mutual fund has lost value as CD rates have gone up. They really don't understand that. We all understand that of course. A number of parameters affect value behavior. The discount rates used, we could also think of these as expected opportunity rates. Also consider the time to final payment, and its cousin, the average time to payment. We'll be talking more about whether the array is front loaded or back loaded and the market's perception of default risk for assets. And finally, I'll be saying quite a bit more about embedded options as I go along, but here I'm talking about calls, puts, extensions, accelerated pre-payments and other kinds of options that are commonly found in fixed income investments.

Let's review quickly the concept of duration. The phrase duration was invented way back in the 1930s by a fellow by the name of Macaulay who defined duration in terms of the measure of longness, a time measure, and it was simply the weighted average elapsed time from now or the point of reference to the occurrence of the cash flows, with the present values of those cash flows used as the weighted factors. Well, Macaulay duration is more interesting than useful. It has to be modified to become useful in the sense that one can use it in investment management and surplus management. The reason for that is that coupons on bonds and cash flows in the liability set are usually not continuous. They're never continuous in bonds. The closest they get are things like mortgage instruments where they're monthly and where the modification is fairly small. But in most cases, with some annual coupon bonds, the modification isn't so small and is important to be done. In mathematical terms, duration is the first derivative of price with respect to yield. Duration is also a linear concept, as opposed to the price yield response curve, which is curvilinear. It could also be described as the slope coefficient which is tangent to the price yield curve at the yield level of our interest.

A related concept, one that may be less familiar to some of us and one that's come much more into prominence in the last several years, is the convexity concept. Convexity is a measure of the tendency of duration itself to change as yield changes. Convexity is the second derivative of price with respect to yield. Graph 1 will illustrate this better than I've done so far with words. This is a picture of the price yield response of a five-year 10% coupon Treasury bond to rate changes. Rates change upwards and downwards by 500 basis points and these look like virtually identical lines, but simply speaking, the lower of these two lines is a straight line and the upper is a slightly curved line. The point here is that in a five-year duration instrument, duration is a very good predictor of value change, even when the value change is fairly large.

By way of contrast, Graph 2 is a 30-year Treasury bond. Again, the bottom line is a straight line. It's pretty easy to see that, here, the upper line is the price yield response line. Now the curvilinear pattern is much more extreme, and the difference in the ability of duration to predict price change in a bond like this is much less than it was in the five-year bond. The duration line is tangent to the curve at the 10% point; at least that's the idea. Were I to draw another tangent line let's say at the 6% or the 14% point, you'd see that the slope of this duration line would be much different. In other words, duration changes a lot in a bond like this. That's what convexity is after all. That leads us to principle number three. Duration explains a large portion of the yield price relationship. Principle number four is convexity explains most but not all of that unexplained by duration. A measure of the immateriality of these other things, the third, fourth, fifth derivatives, is that nobody has given them a name yet, so they really aren't very important, even in bonds like the 30-year (Graph 2). Principle number five says that, when duration is relatively short, less than five years, or when the yield change is relatively small, duration is a fairly exact estimator of value change. Under these conditions, convexity is relatively unimportant.

Let's shift gears and look at some familiar liability arrays and their behavior (Graph 3). This is a picture of the retired lives liability set. It comes from a real situation, not at all untypical, the retired lives of a pension plan. Again, the bottom line is the duration line and the upper is the price yield or value yield response curve. You can see that this graph is similar, but not as extreme as was the case with the 30-year Treasury. This particular set of liabilities has a modified duration of about 6.25 years. Of somewhat more interest would be to look at all of the vested liabilities of this very same pension plan (Graph 4). That would be the retired and all vested, terminated, disabled and so forth. As before, the bottom line is duration and the upper line is the value yield response curve. This situation has a modified duration of about 8.29 years. Now, let's compare this curve, if you will, to a bond that has the same duration, and I've selected this bond so that it will have the same duration. It's an 18-year Treasury bond and has essentially

GRAPH 1

TREASURY BOND – 5-YEAR MATURITY

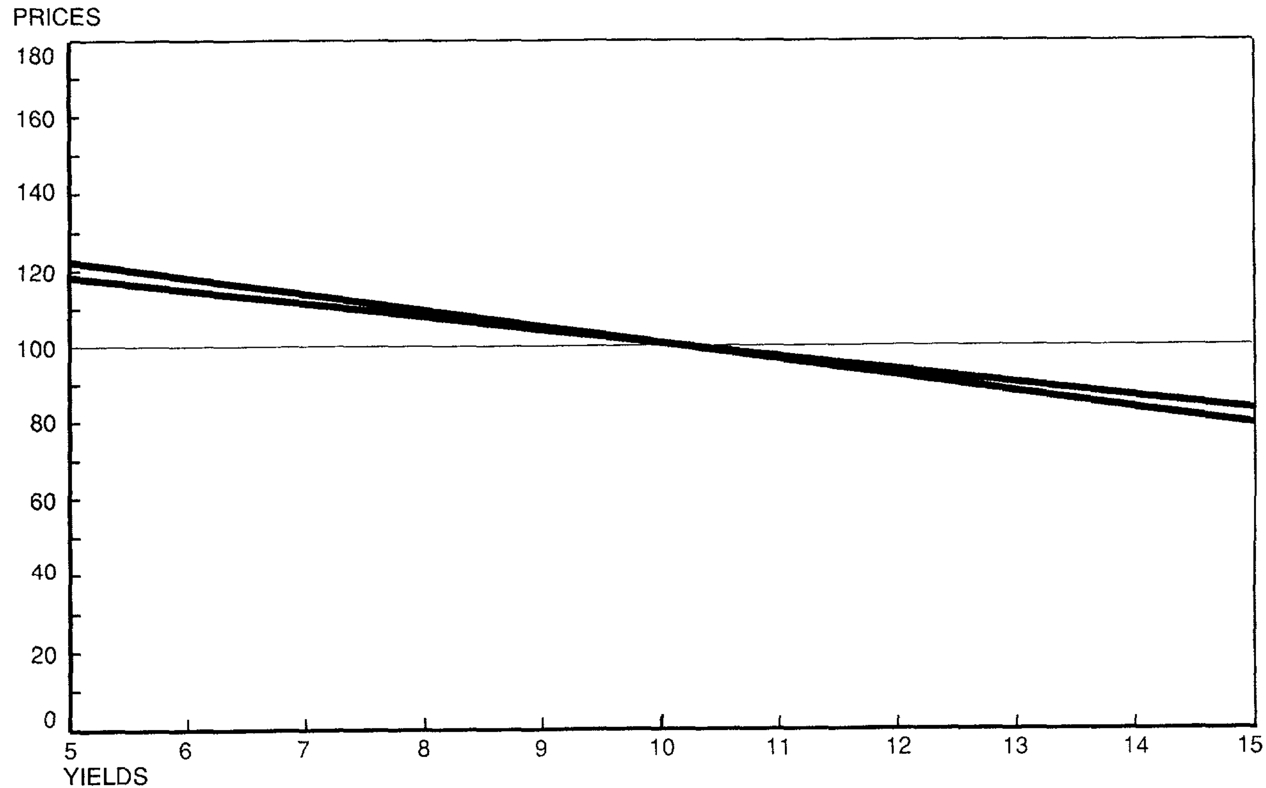

GRAPH 2

TREASURY BOND – 30-YEAR MATURITY

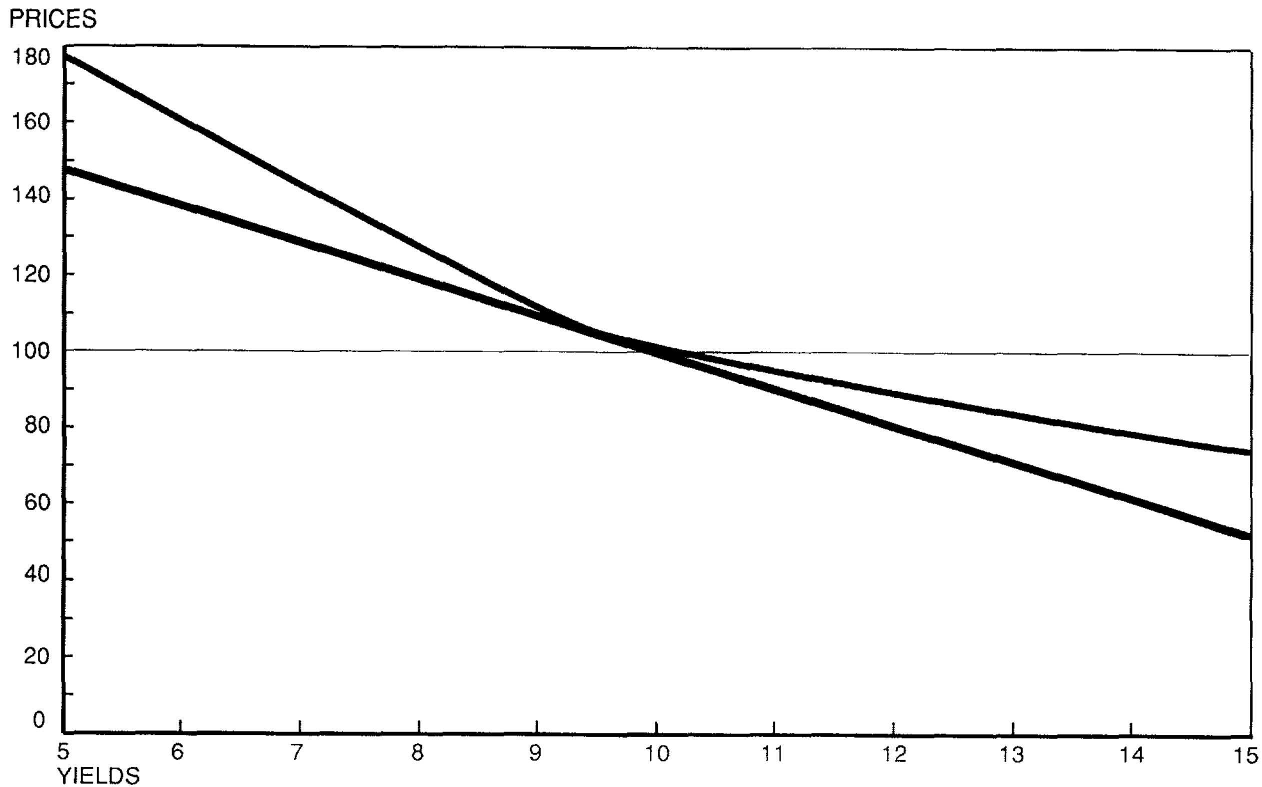

GRAPH 3

RETIRED LIVES LIABILITY

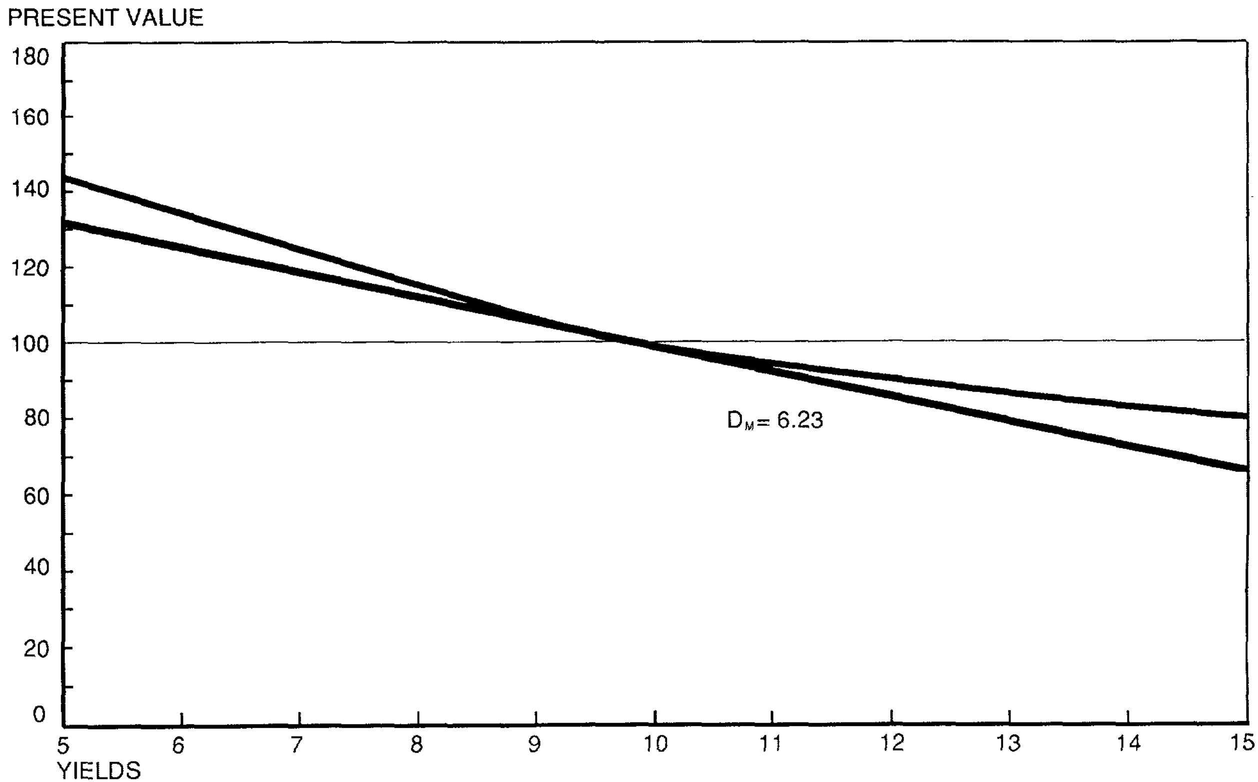

PANEL DISCUSSION

GRAPH 4

ALL VESTED LIABILITIES

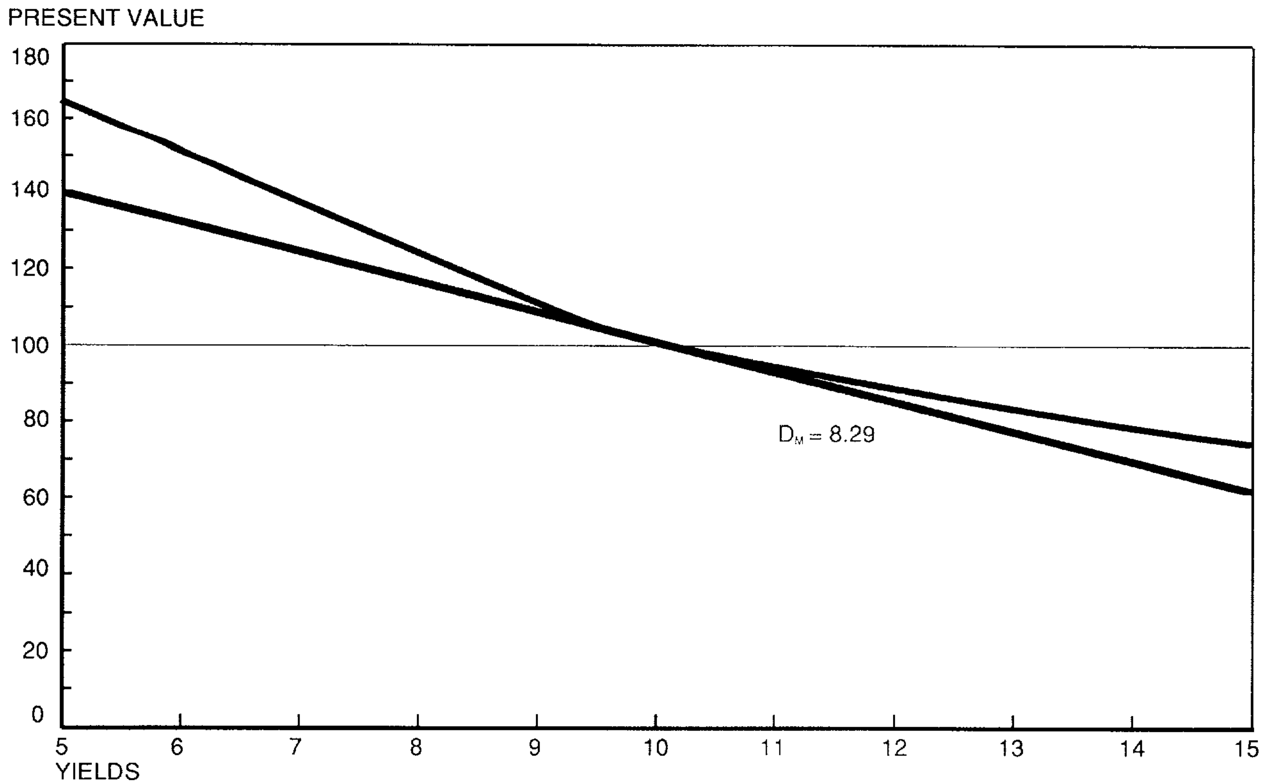

the same duration as the liability set that we just looked at. On Graph 5, you'll have to imagine that I'm putting these two curves together. These are those two situations, the 18-year Treasury bond and all of the vested liabilities of the pension plan that have identical durations. You'll notice that their price value to yield response curves are not identical. They're a little different, especially at the tails. Again, the duration slope, which is a straight line, describes both of them. To magnify that difference, Graph 6 describes the value difference in dollars per $100 of asset or liability value when rates were at 10%. In this case, when rates dropped from ten to five, the assets increase in value $7.00 less than the liabilities increase in value. The liabilities increase more than the assets do. What this is saying is that even a perfectly duration matched set of assets and liabilities doesn't necessarily produce an immunized result in the sense of value changes.

To summarize this, we use principle number six. Cash flow arrays of identical durations may not have identical behavioral characteristics. If the durations are identical and the convexities are not, the behavior will differ, and that's true even when there aren't any embedded options. This is true even in the unlikely situation that yield curve shifts are parallel, that is, equal magnitude at all points on the yield curve. Another principle would be cash flow arrays that are more even have more convexity per unit of duration. That certainly is the case in comparing a set of liabilities connected with a pension plan to bonds with bullet maturities, such as a Treasury bond. A retired lives liability array then has more convexity, typically, than a portfolio of bullet bonds with matching duration. A special case that would be an exception to that principle would be what some folks call a dedicated portfolio, a cash-matched portfolio in which the bonds selected are selected so that the cash flows of principal and interest exactly match the liabilities. In that case, duration and convexity and cash flow are all perfectly matched, and there are no problems. The problem with that solution is that it tends to be a somewhat more expensive solution in terms of implementation, usually by anywhere from 1 to 4% in an immunized portfolio, and it has some other negatives as well.

Let's turn now to the problem of embedded options. There are three ways that people tend to deal with embedded options in bonds. The first is to ignore them, pretend they aren't there. This always gives the wrong answer. It always gives the wrong duration, convexity and, in fact, yield answer, and sometimes the error is quite large. This choice of the three choices is still, surprisingly, in very widespread use. In fact, five to ten years ago it was virtually in universal use. Nobody paid attention to options. Another reason, I think, was that five or ten years ago people had gone an entire career and never seen interest rates go down. Their options were more a figment of imagination than anything real. Now we've seen that rates can come down, and that can be harmful.

The second concept, which is in fairly widespread use, is connected to the idea of yield-to-worst, and many of us, I guess, know what yield-to-worst is. One calculates the yield to maturity on a bond, then one calculates the yield-to-call using the call price and the first call date, and the lower of those two yields is the yield-to-worst. That concept can be extended into the area of duration. If the bond is priced on a yield-to-call basis, then the inference is that it is thought that the call will occur with a probability of one and the duration is based then on the call parameters. So if the bond is callable in two years, then the duration will be nearly two years. If the bond is callable tomorrow, it will have a duration of very near zero. If, on the other hand, the yield-to-worst is equal to the yield-to-maturity, then the bond is assumed not to be called with a probability of one; that is to say, no probability of call and yield-to-maturity and time-to-maturity determines the duration. This approach does occasionally give essentially the right answer. When the probability of call is so low, that is when rates are very, very high or when the probability of call is very, very high when rates are quite low, then you'll get the right answer. The easy way to think about this is a light switch. In our dining room we have a rheostat which can adjust the light intensity infinitely. In most other rooms we have a light switch. It's either on or off. This second option is the light switch option.

The third option, which I'll be talking about in a moment, is the rheostat. The third approach which is coming into fairly widespread use among Wall Street bond houses and, I think, in the actuarial community, is an approach to the options that deals with the characteristics of the bonds, the behavior of those bonds, and understanding the option has a value which changes with interest rates. The problem here is that this is not mathematically simple or concise. It's inappropriate to do it using, for example, the Black Scholes model, which is not simple in concept, but it's simple to apply. Black Scholes is a closed-form solution. It doesn't work for bonds for a couple of reasons. In fact, it wouldn't work for stocks where the option runs for a very long time,

GRAPH 5

DIFFERENTIAL CONVEXITY
Vested liability vs. UST-18

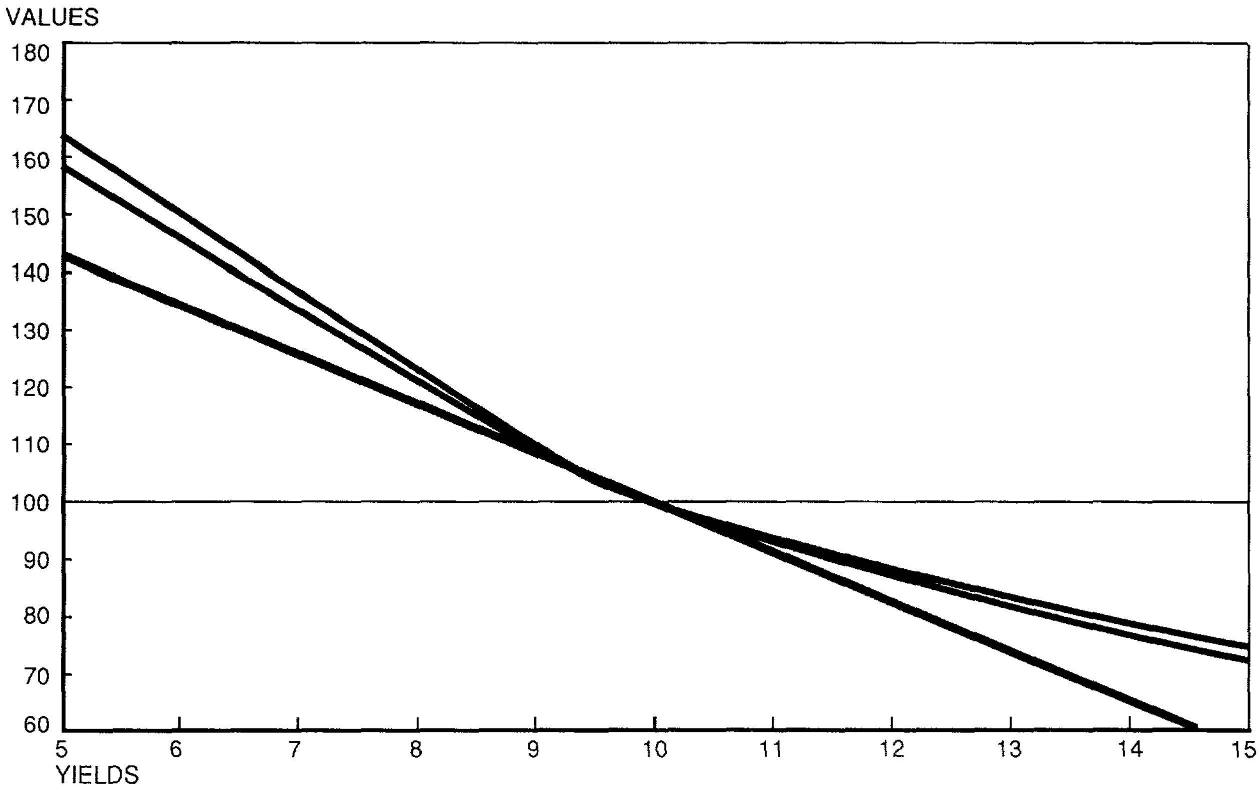

GRAPH 6

EFFECT OF DIFFERENTIAL CONVEXITY

Durations equal and bonds noncallable

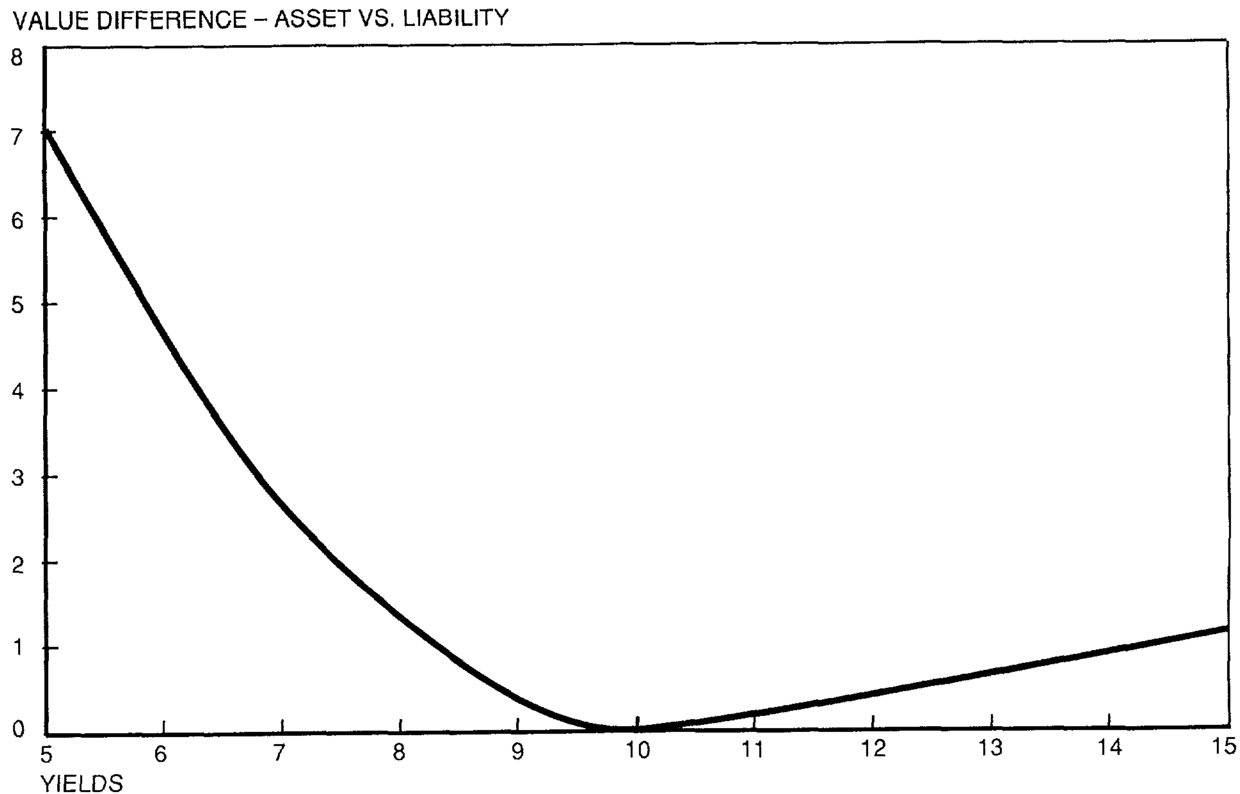

which is the case usually with embedded bond options. The easiest way to think about this last approach is a simple algebraic equation. When I buy a callable bond, what I'm really buying is a non-callable bond and I'm selling an option, so that the callable bond is the sum of my non-callable bond that I own and the option that I sold. That's the value equation. Similarly, the duration of the callable bond can be thought of as being the summation of the duration of the non-callable bond and the duration of the option, and likewise with the convexity. The difficulty here, and a fine risk concept, is that for some of these values it's not possible with precision to know what these values are. The approach that's usually used is to put a value on the option, and now we're back into this difficult problem of valuing the option by processes that are not mathematically simple and, in fact, are computer pigs. They require lots of computer power to do them well. Non-callable coupon-paying bonds, in the current market, have modified durations on the order of zero to eleven. Options can have very large sensitivities to yield changes; that is to say very large durations and very high convexity. Referring to the formulas of value, if the value of the option is very small, then its impact on duration and convexity is also very small, even though the value of its duration and convexity are large. The value is small in terms of its weight in the first of these formulas. As the option value increases, its impact on the behavior of the bond to which it's attached becomes much more important. And finally, its value-weighted duration can equal the value-weighted duration of the non-callable bond, and you get a condition in which price doesn't change regardless of how much rates change.

For example, imagine a 16% coupon bond that's callable a week from today at 105. The market's going to pay very little attention to changes in rates in valuing this bond. It's going to be valued at around 105 or maybe a smidge higher, irrespective of interest rates. It just doesn't matter. What it means is that the duration is very short and the reason is that the net duration, the net of the two pieces, is equal to zero. Now also the option's value-weighted convexity can exceed that of the bond. In other words, unlike duration where it can't become greater than that of the bond, it can actually become larger and the aggregate convexity can become negative.

The top line on Graph 7 is an 11% 22-year bond, non-callable for life. The lower line is an 11% 22-year bond that's callable currently at 107. What you're seeing is that it doesn't gain nearly as much value as its non-callable brother. The difference in market values can be extreme. These are two bonds that had virtually equal values when rates were high. But the non-callable bond, in a 6% environment here, is worth $40.00 more, or 40% more than its callable brother. Here you see this effect in embedded options when rates come down. Still another way to think about this is what's the behavior of the duration itself as rates change? The upper line then is the non-callable bond, and you can see that duration is ever-increasing at all points. What that means is that duration has gone in the direction one wants it to go. It's this increase in duration that's causing the positive and therefore desirable convexity. By contrast, the callable bond has an increase in duration only at rates above about 14%, and at any rates thereafter, any change is adverse and it's that condition that causes negative convexity, which is harmful. This brings us to principle number seven. True immunization then is only possible when both duration matching and convexity matching are present. Any deviation from this double match-up condition places the insurance company at risk. The question I pose here is, and for the pricing actuary in particular, is the risk induced by absence of this double match-up compensated for with a risk premium in pricing insurance and annuity products? Remember, options always have a value greater than zero. In fact, the option's value largely derives from its powerful, positive convexity. Borrowers, that is bond issuers, pay for call options in the form of lower issuance proceeds per unit of interest expense, so callable bonds have higher yields than if they weren't callable. If these higher yields, which have been paid to the lender, the insurance company, by the borrower, have been passed along to the customer in the form of higher guarantee rates, then the insurance carrier has sold convexity and given away the proceeds of the sale in the pricing process. This type of pricing, when it happens, represents a major bet on interest rates. Over a very long period it says, I'm willing to bet that interest rates will remain at or near or higher than current levels for the next 20 years.

Quite apart from these questions about convexity and embedded options, there are other factors to consider in thinking about behavioral characteristics. Duration and convexity calculations are always based on the assumption, almost always anyway, on a flat yield curve and that, when changes occur in that curve, they will be parallel. That assumption almost never is true in practice. And this assumption that is usually made, but which is almost never true, has particularly important implications when the incidence of asset and liability cash flows is materially different, that is, when they don't match up. It's not a cash matching situation, even when

GRAPH 7

COMPARITIVE MARKET VALUATIONS

Callable vs. noncallable bonds

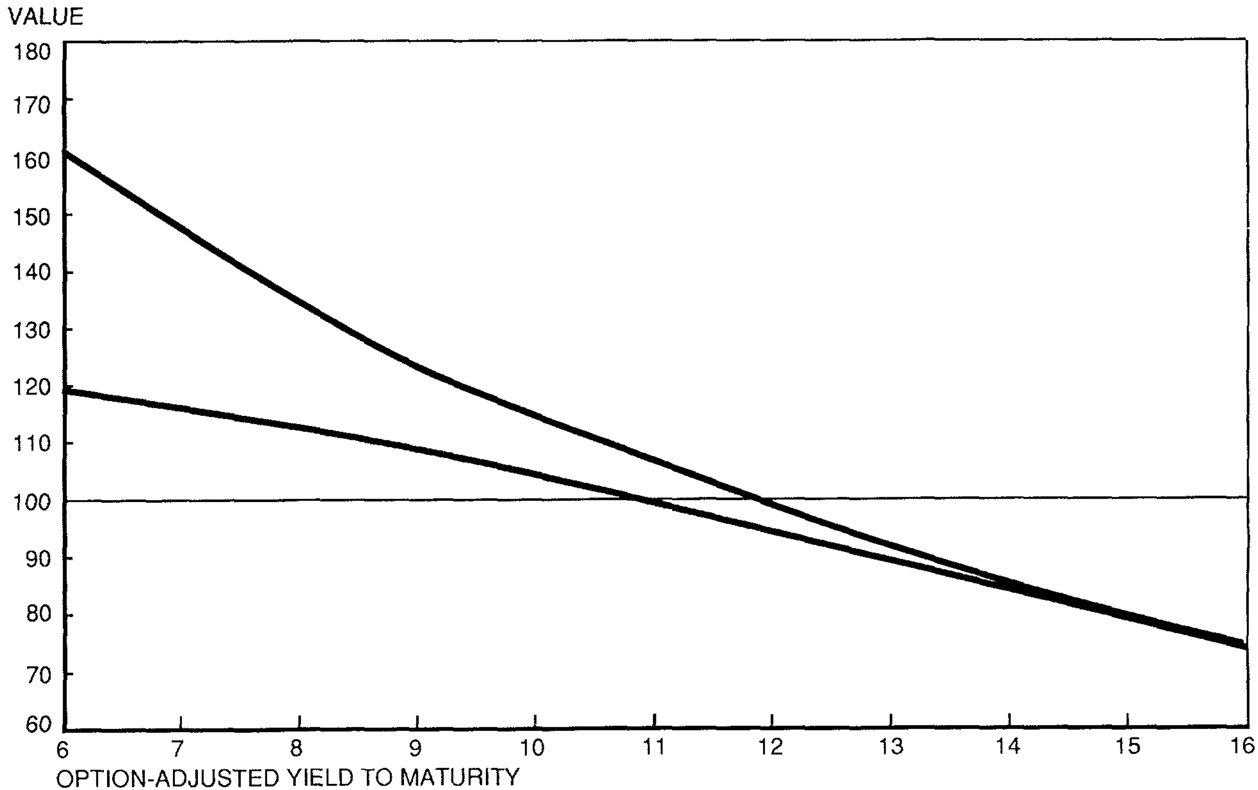

durations are matched. In such situations, an adverse change in the shape of the yield curve can have a profound effect. That would be particularly true when the yield curve has been inverted and is still, to some extent, inverted, but really flat. When it snaps down to its more normal shape, that can have an adverse affect on an asset/liability situation in which the cash incidence is not matched, in which the convexity is not matched.

Finally, as a kind of summary statement about the things that affect the value of bonds. What I've done is to try to identify all the elements that cause bonds to have greater or lesser value: anticipated high inflation, anticipated inflation of any kind for that matter, illiquidity, more liquidity or less liquidity than Treasury bonds, for example, any risk of default relative to Treasuries. All of these are bad things. The one thing that's good in a bond is positive convexity. The market will pay more to get positive convexity.

MR. D. ANDREW HALL III: As you probably know, but don't necessarily think about all the time, an insurance company's business is to write options on both sides of the balance sheet. That worked out great for a long time when interest rate volatility was low, because when you write options, you make money if volatility is low. It's been a lot trickier in the 1980s since volatility picked up. I'm going to be talking about the options in the asset side of the balance sheet, how you evaluate these assets, and determine their price sensitivities net of all of their embedded options.

First of all, I have a list, by no means complete, of assets that do have embedded options. Corporate bonds have embedded options, the call options that Bob was talking about. There are also puttable bonds. In that case, you are long the options so you get a lower stated yield. With callable bonds, as Bob said, you get a higher stated yield than with a non-callable bond. There are things called indexed sinking fund debentures. There are all sorts of options embedded in corporate bonds. Fixed-rate mortgage pass-throughs have embedded options. The options there are pretty complicated because there's this whole portfolio of options in a mortgage-backed security (MBS). Every borrower inside a mortgage pool has an option to prepay in any month. Then the markets have come up with things called IOs and POs, which are stripped MBSs that leverage the prepayment option. Then there's collateralized mortgage obligations, CMOs. There you do all sorts of strange things with the prepayment option. You can allocate it many ways: You can leverage it, or you can take it away from part of the CMO, which leaves the prepayment option stronger in the other parts of the deal. Finally there are adjustable rate mortgage securities which have a plethora of embedded options.

In Graph 8, I'm going to talk about two ways of thinking about a callable bond. Now, the simple example I have here is a ten-year no-call five. In other words, it can be called in five years. Now, if that bond is trading at a discount to the market, which means that its coupon is lower than where the corporation could currently issue a bond, you will see that it's most likely not going to be called. In that situation you can think of yourself as being long a ten-year non-call bond and short a five-year call option on a five-year bond, because if rates go down, the corporation will exercise its call option and take away the last five years of your then premium security. Equivalently, if the bond is trading at a premium, and you expect it to be called, you can think of yourself as being long a five-year non-callable bond and short a five-year put option. If you take the difference between the two sets of boxes, a thing that we do in mathematics a lot, you actually get an equation that is known as put/call parity. It's a very important concept in option pricing theory, but we can talk about that later. Similarly, a puttable bond can be thought of in two different ways as well. With a puttable bond (Graph 9), the investor is long the embedded option. The issuing corporation is short that option, so they can issue the bond at a lower initial yield. If you're long a puttable bond, you can think of yourself as being long a ten-year no call bond and long a five-year put option. That's the way you think about it, if it is trading at a premium. You're not liable to put it back to them if it's trading at premium currently. Similarly, you can think about being long the puttable bond as being long a five-year non-call bond and long a five-year call option. These relationships show up the symmetry between puts and calls.

Now, I'm not going to be talking in detail about how to evaluate the options in corporate. I'll be devoting the time on option pricing to mortgages, but one way you can think about getting at what the real underlying yield is, what your give-up is if you buy a callable bond, is in Graph 10. We turned around one of the first equations, and you can think of the price of the non-callable bond as being equal to the price of the callable bond plus the price of the call option. The simple example here is a real bond a couple years ago. Suppose you buy a 30-year 12% coupon bond that's

GRAPH 8

Callable Bond
10-Year/Callable in 5 Years

Long		Long		Short	
Callable Bond	=	10-Year Non-Callable Bond	−	5-Year Call Option	(Discount)

-- or --

Long		Long		Short	
Callable Bond	=	5-Year Non-Callable Bond	−	5-Year Put Option	(Premium)

Puttable Bond
10-Year/Puttable in 5 Years

Long		Long		Long	
Puttable Bond	=	10-Year Non-Callable Bond	+	5-Year Put Option	(Premium)

-- or --

Long		Long		Long	
Puttable Bond	=	5-Year Non-Callable Bond	+	5-Year Call Option	(Discount)

GRAPH 9

GRAPH 10

Calculation of Call-Adjusted Yield

Price of Non-Callable Bond	=	Price of Callable Bond	+	Price of Call Option

Assumptions

30-year, 12% coupon bond, callable after 5 years at 110
Bond is priced at par to yield 12%
Option premium is estimated to be $1.89

Calculation

Price of Callable Bond	$100.00	
+		
Option Premium	1.89	
Price of Non-Callable Bond	$101.89	⇒ Call-Adjusted Yield = 11.77%

callable in five years at 110. The bond is priced at par and yields 12%. The value of the option, using a fairly constructed option pricing model, is estimated to be $1.89 for every $100 of principal. So if you look back at the top box, that means that the underlying non-callable bond is worth about $101.89. You then calculate its yield with the 12% coupon, and it comes out to 11.77%. In other words, you're paying, essentially, 23 basis points for that call option.

Now I'm going to talk a little bit, and this goes along with what Bob was talking about, about the durations of a callable bond and a puttable bond (Graph 11). What we have on the left side is the duration of a callable bond, a ten-year no-call bond and a five-year no-call bond. The top line there is the duration of the ten-year bond. It's a function of interest rates which are on the lower axis. The lower line is the five-year bond. If market interest is well below your coupon rate, you can assume the corporate treasurer will call that bond, and so the callable bond trades just like the five-year bond. That is, with probability close to one, the bond will be called. Similarly, if you have very high interest rates, you're going to assume that it will not be called because the corporation would have to issue debt at a higher coupon, so the bond trades just like the ten-year bond. In between, where it switches from the track in the top curve to the bottom curve, is where the negative convexity comes in strong. That's where the option is close to the money. You don't know if it's going to be called or not. It depends on what happens to interest rates. As interest rates rise, your duration goes up sharply. That's not good. You don't want that to happen because as interest rates are rising, you're losing value in your security. If the duration is increasing, well, then you're losing value at an accelerating rate, and that's what negative convexity is all about. The other graph is very similar, except now it is a better situation. You're long the option, so you have positive convexity. At very low interest rates, the puttable bond trades like the ten-year bond. The investor is not liable to put that bond back to the corporate treasurer if he is getting a higher than market coupon. For a very high rate, he is liable to put it back so he can get his money back at par and put it to work at a higher rate. In between, as rates move up from below the strike yield to above the strike yield, your duration drops sharply as interest rates are rising. That's good, you want that. That's strong positive convexity.

Now I'm going to move into MBSs. What do those look like? Well, using the box in Graph 12, if you're long an MBS, that is the same as being long a 30-year non-callable amortizing bond, just a straight non-prepayable mortgage, and short N, where N is the number of mortgagors in the pool, times 360 call options. A complicating factor here, if that's not complicated enough, is that those call options are not exercised efficiently; for example, some people don't prepay 16% mortgages, while some people do prepay 7% mortgages. In option pricing theory you call this behavior irrational exercise. Now, this is where we come into actuarial work. This is what I brought as an actuary to mortgage research; prepayment modeling is the same mathematical problem as mortality. It's just a little more complicated. Mortality depends on age and health and a few such parameters, whereas mortgage prepayments have more determinants. In fact, the most applicable number we use to categorize prepayment rates is called SMM, which stands for Single Monthly Mortality. It's a measure of how much a mortgage pool prepays in a given month.

What are the determinants of fixed rate mortgage prepayments? Well, the first and most obvious one is the difference between what the mortgagor is paying, his coupon, and what current mortgage rates are. Clearly if you're paying 14% and current mortgage rates are 9%, you have a strong financial incentive to prepay. The second one is the age of the mortgage. You very rarely have a strong incentive to prepay if you just took out a loan and you paid the issuer two points. It would take some extraordinary circumstances for you to prepay your loan quickly. Another strong determinant of mortgage prepayments is actually just the month of the year. People move much more often in May, June and July than they do in December, January, and February and that turns out to be an important determinant in mortgage prepayments. Then there's the rate history over the life of the loan. If you have a 12% mortgage issued a year ago, for example, and mortgage rates come down to 9% very suddenly, you'll see that pool prepay extremely quickly for a while, as all the interest-sensitive borrowers in that pool prepay. They say, hey, this is in the money. I've got a nice option here. I'll prepay it and refinance my loan at lower cost. Suppose rates go back up to 12%? They no longer have any economic incentive to prepay. If rates come back down again, the response the second time around is measurably lower, and this is the sort of thing that leads to the 16% mortgages that some of you may be lucky enough to have in your portfolio, prepaying at a very low rate of 20% a year and therefore having extremely high yields. This phenomenon is called prepayment burnout. This leads into the irrational borrower behavior. It seems as if anybody with a 16% mortgage who doesn't prepay is behaving irrationally. That

GRAPH 11

Effects of Options on Corporate Bond Price Sensitivity - Duration

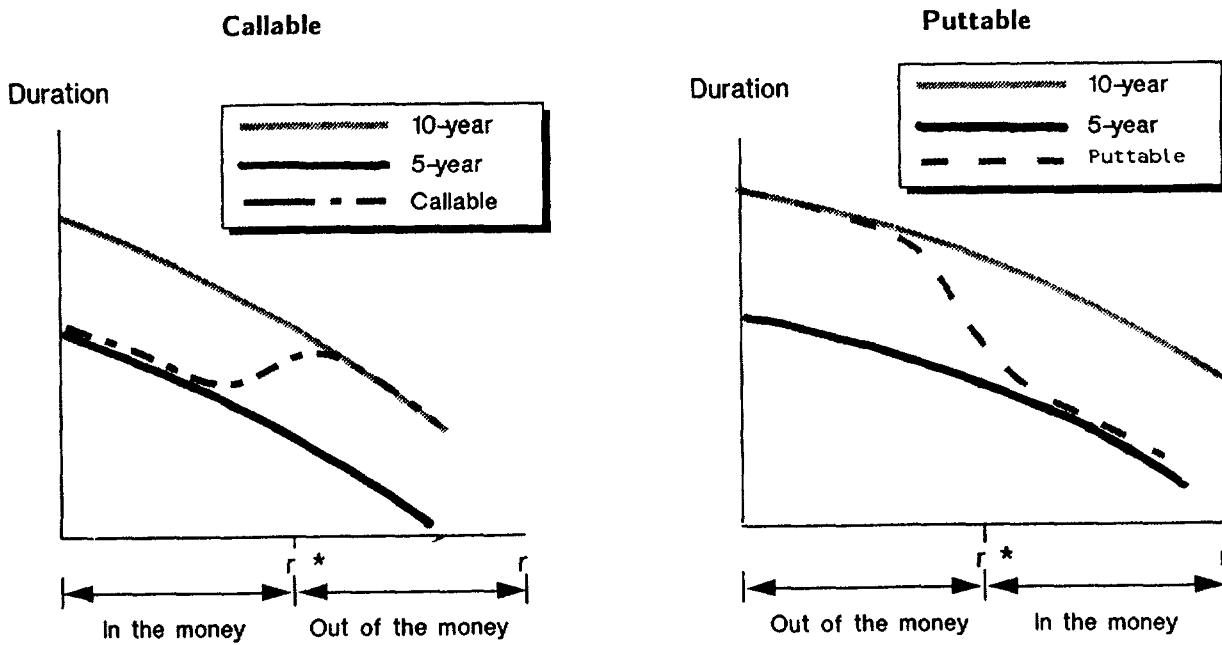

Mortgage-Backed Securities (MBSs)

Long		Long		Short
MBS	=	30-Year Non-Callable Amortizing Bond	–	N* x 360 Call Options

* N = Number of loans in pool

may not be true. He may not be able to qualify for a new loan. He might have been in Houston and only made a 10% down payment and watched the value of his house go down by 20%, so he owes more on the house than the house is worth, or he might want to stay in the house and be able to scrape enough money to pay that mortgage, but not be able to qualify for one. Similarly, it seems irrational for somebody with a 7% mortgage to prepay, but they might have to move or they might have died. There are logical reasons for those people to do that. The last prepayment determinant, which is quite hard to model, is the prevailing economic conditions, what borrower sentiment is. People think rates are going down. They think they're going up. All sorts of stuff like that. As I say, prepayment modeling is a mortality problem, and in fact, our mortgage valuation models have multi-decrement mortality tables as the basis of the prepayment models.

I'd like to quickly go through the history of the MBS valuation, how people worked their way up to actually pricing the options properly. Initially, from day one up until somewhere around 1985, people priced mortgages on the basis of a 12-year prepaid life. What that meant was that they assumed there were no prepayments at all for 12 years and then the whole pool of loans prepaid at once. Well, as you can see, that would grossly overvalue a premium. People do prepay premiums sooner than that, so if you've got a premium there at a certain 12-year prepaid yield, you would not have achieved that yield. Conversely, it grossly undervalued discounts. If you're assuming no prepayments on this security for 12 years, when in fact the lowest we've ever seen is 3-4% prepayments a year and that's very rare; 6-9% on the deep discounts is common, and you purchase the loan at a given yield, then you achieved a far higher yield as the discount MBS prepaid during the 12 years. A lot of insurance companies realized that and made a fortune by purchasing discounts on a 12-year prepaid life basis. The second way, when people started to get a little more sophisticated, they used Federal Housing Administration (FHA) experience to estimate prepayments. The FHA experience is based on five-year average prepayment rates. They looked over five years and how likely it was to prepay the first month, how likely it was to prepay the second month, how likely it was in each month over that whole five-year period; exactly the way you model mortality. For mortality such a study is appropriate. It's not appropriate in mortgages because the prepayments depend on interest rates as well as age. The market interest rates were different for a mortgage originated on the first month of the five-year period and one issued in the middle. Now, if you're just trying to estimate prepayment rates on age only, and those are different interest rate environments, you're losing a whole lot of information that is critical for evaluating a mortgage. The next step in the valuation history, and one that's still in current use and I believe is faulted, is assuming a constant prepayment assumption. Those are measured two ways that are equivalent, CPR which is an annualized rate, and SMM, which is a monthly rate.

As I said, SMM stands for single monthly mortality. If you assume a constant rate for the life of the loan, you ignore the aging, the tendency of a mortgage to prepay very slowly at first and come up to a level and then after ten or 15 years of life increase again as people start to get older and demographic factors set in. It also ignores the fact that interest rates change. These cash flows aren't fixed; if interest rates change, the prepayment rates will change, and therefore, it gives you extremely poor duration estimates for anything except a very deep discount or a very high premium. They used modified Macaulay duration, which is not a measure of price sensitivity in an interest-sensitive instrument. It's just not. Finally, in 1987, people went live with option-adjusted spread models. A little later I'll go into a lot of detail on a particular Old Age Security Act (OAS) model, the one we use to evaluate adjustable rate mortgages (ARMs). We have proven that OAS modeling has historically been an excellent indicator of value and that it gives accurate and unbiased estimates of duration and convexity. I'll conclude the talk with some examples of these historical tests.

Now, very quickly we have stripped MBSs, a derivative product of MBSs. The two big categories are IOs and POs. Now, if you get the interest-only cash flows, that means exactly what it says. You get just the interest off of the outstanding principal, so what you want is the principal to be outstanding a long time so you get that interest for a long time. If the mortgage completely prepays, if the pool has 100% prepayment next month, you don't get anything. You paid some money and you get zip. You get a negative yield. So you benefit very well if prepayments slow down below what you initially expect. In other words, interest rates go up, prepayments slow down, and your value increases, so IOs often have negative duration; not always, but often, because they have an inverse relationship between price and yield. The principal only is just the converse. You do extremely well if prepayments speed up. You only have to pay $40 or $50 for the right to a $100 in principal, so if it completely prepays tomorrow, you've doubled your money and you're very happy. So you benefit if prepayments speed up.

The second category of derivative MBS debt is CMOs. By no means do we have time to describe any of these very carefully at all, but there are a lot of different classes. The whole idea is carving up that prepayment option and giving different pieces of it to different people. The planned amortization classes (PACs) are similar to corporate's. You're trying to target a prepayment range from quite slow to quite fast. You can maintain the cash flows on this bond over this protected range. You've taken away some of that prepayment option from this class, but there's conservation of dollars here. You can't just take average life volatility away there and not add some more volatility someplace else. The non-PAC bonds, in a CMO that has a PAC tranche, become much more volatile, and I'll be happy to talk about that in more detail, time permitting. Then there are floating rate CMOs. There are super POs which incredibly leverage the prepayment option and make a lot of sense for certain insurance portfolios. So, these are different ways that we, on Wall Street, have carved up mortgages. They're poorly enough understood that if you use careful option pricing techniques and the guys down the hall are just using static cash flow analysis with single monthly mortality (SMMs) and conditional prepayment rates (CPRs) and public securities association rates (PSAs) to some fixed cash flow thing, you can often pick them off and find bonds that offer good value. If you don't do it correctly, you can end up getting a tranche that has very poor value.

Now we come to ARMs, and these have an enormous amount of options built into them. First of all, since they're adjustable rate, the coupon on them changes periodically, and that is, in fact, an option itself. The periodic rate cap is an option. A large number of the ARMs in the market have coupons that are only allowed to change so much: typically 200 basis points a year. All ARMs in the market now have lifetime rate caps. The coupon cannot exceed a certain amount fixed at the time of issuance. Then there are a number of ARMs that have payment caps, where the payment can only change by so much per pay reset. If your payment is so restricted that it doesn't go up high enough to amortize the loan, ARMs add the difference of what the borrower owes to principal, so they can have negative amortization, i.e., the balance can go up. Conversely, you can have accelerated amortization as well. There's maturity extension in certain ARMs, which means that the borrower, when faced with a payment increase, has the option of extending the maturity and therefore having a lower payment rate. It's not as common as it used to be. Very common now in the market are convertible ARMs, which means that the borrower has the option, typically in between the first and the fifth year, of converting his ARM into a fixed rate mortgage. That doesn't mean if you invest in ARMs you're suddenly going to own a bunch of fixed rate mortgages. What happens to the security holder is he gets prepaid at par, so it looks like a prepayment, but it does affect the cash flows that you get. Finally, starred at the bottom, you have the classic prepayment option.

Because ARMs have the most embedded options, I will talk about the way that we evaluate ARMs net of their embedded options at Morgan Stanley. I use a flow chart that applies to ARMs (Graph 13). This is the same mathematical technique that we use to evaluate the options in corporate bonds, fixed rate mortgages, and CMOs, but ARMs have a few more bells and whistles. What we're trying to do here, and this is sort of the guts of the talk, is show how to evaluate some complex security net of its embedded options, and we're trying to calculate what we call the option-adjusted spread. By definition, the option-adjusted spread is the spread between this complex security and the Treasury yield curve. We use the Treasury yield curve as the basis across the board at Morgan Stanley. So the first input to the model, up in the upper left, is the current coupon Treasury yield curve. This model is a simulation model, so using various rules that I'll talk about, we create hundreds of interest rate scenarios based on that initial yield curve and volatility assumptions. Each path in the simulation goes out for the life of the security you're trying to evaluate and rates change along the path periodically. In our fixed rate model and our CMO model, we just carry two points on the Treasury yield curve, the 90-day Treasury Bill rate and the ten-year note rate. That's all you need to evaluate a fixed rate mortgage. To evaluate an ARM, we need a lot more. ARMs are tied to a number of different indices, so we carry along the one-month, the three month, the six month, the one-year, three-year, five-year and ten-year. It's a huge model. It takes up about 50 mega bytes on the computer and dims the lights in New Jersey when I turn it on. Each point on the yield curve is generated with its own volatility. We don't assume the same volatility for a one-month rate as a ten-year rate. Those have very different historical volatilities. Short-term rates have had higher rate volatility than long-term rates. Since we're pricing a bunch of long-dated options, all the options embedded in a 30-year mortgage in this example, we use a five-year historical volatility assumption. In other words, we look back at daily yields on all these Treasuries for five years, calculate what the actual rate volatility was and that's what we put in the model. We don't know if you should use three-year, four-year, five-year,

GRAPH 13

Basic Principles of Valuation

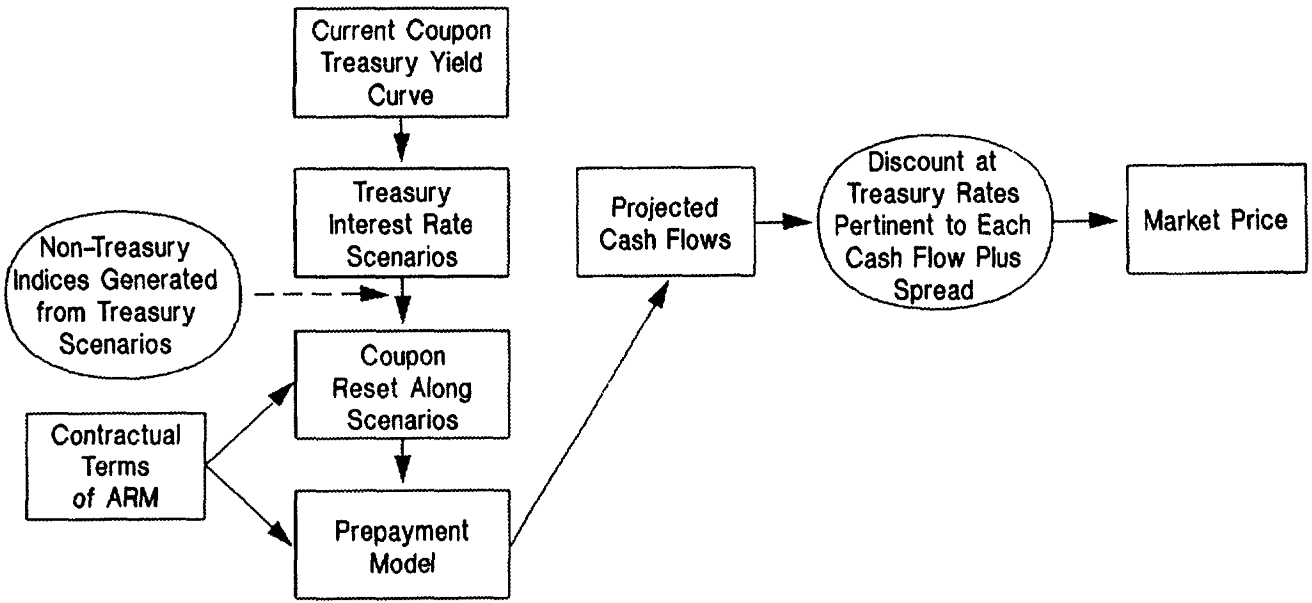

or ten-year volatilities. We just know it should be long-dated and we have demonstrated that, so we, by convention, use a five-year volatility rate. We also use a correlation matrix between the different rates on the yield curve. That allows the yield curve to steepen, flatten or invert. These correlations are high, but they're not 100%. The yield curve does change shape. This is very important when you're evaluating mortgages because borrowers look at both the long end of the yield curve for fixed rate mortgages and the short end for ARMs. They might not know that's what they're doing, but they are. So it's very important when you're modeling prepayments to have both ends of the curve.

As Bob had said, the Black Scholes option pricing model doesn't work for bonds. Neither does purely random generation of these interest rates. You wouldn't price Treasuries properly. We're trying to compare some complex security to Treasuries, so the first requirement we put on our model is that it prices Treasuries properly. Remember, we have all these hundreds of interest rate paths. Take any given Treasury from that current coupon Treasury yield curve, put it in its fixed cash flows for every path. There's no option in there. They're going to be the same cash flows no matter what happens to interest rates. Then discount them, using the rules we'll talk about later, along each path using an option-adjusted spread of zero. The spread between Treasuries and Treasuries is, by definition, zero. You'll get a different present value for each path because you've discounted with different rates. Average those present values. We require that you get exactly the price of that Treasury. A lot of people don't when they use their option-adjusted model. If you randomly generate scenarios, that won't happen. You have to be more careful than that. So you'll be trying to compare some security to Treasuries using a model that can't even price Treasuries. That doesn't make a lot of sense. Another requirement we put into the model is that if you price call and put options on any Treasury in that current coupon Treasury yield curve, that the prices of the put and call will satisfy put/call parity. So these arbitrage free conditions we've put on require that the model prices Treasuries exactly and options on Treasuries consistently.

The next step, whatever security you're evaluating, is the estimation of the security's cash flows along each interest rate path. I'll talk about how we do that with ARMs. I think you can extend it to how you do it with other securities. First of all there are some ARMs that are tied to non--Treasury indices, so we haven't generated enough rates yet. If you thought 50 mega bytes wasn't enough, well, let's try a few more. We use separate models to estimate these non-Treasury indices like the 11th District Cost of Funds or the National Contract Rate or the National Cost of Funds, all of which are indices that ARMs are tied to. We use separate, typically regression, models to estimate these indices as a function of the Treasury rates along each scenario. Then we look at the contractual terms of the ARM, what is its coupon, its margin, periodic cap, lifetime cap, reset frequency etc. and generate the ARMs coupon along each path in the simulation. So now we have hundreds of sets of coupons. Then we pick the appropriate prepayment model; we have different prepayment models for different types of ARMs because they do have different behavior. Using the prepayment function we estimate the prepayments along each path. Between the coupon and the prepayments, we have the cash flows of the ARM. We estimate what the option-adjusted spread is. We don't know what it is. We know what the market price of the bond or the security we're trying to evaluate is. Estimate what that option-adjusted spread is. Go to the last point on the first scenario. Take that last cash flow, a couple of nano-pennies by that point, and add the estimated OAS to the short-term rate that's out at the end of that path, discount that cash flow back one period, add it to the cash flow there, add the estimated OAS to the second to last interest rate on that path, discount it back one more period. This is the critical part of the option pricing model that makes the put/call parity work. It makes you price Treasuries properly. You discount not with a single rate for all the paths, but you discount using the rates along the path that you projected the cash flows along, adding the OAS to the short-term Treasury rate. That's the crucial thing. That's how we got put/call parity back in 1984-1985 when we first started generating options models that priced debt options consistently. You do that all the way back to the beginning, and you get a present value for the security on the first path. Once you have repeated this procedure on each path, you've converted all of these cash flows into hundreds of present values. Average those present values; that is what the price of the security would be, if you have estimated the OAS properly. Compare this average present value to the market price. If it's the same, you are done. Your estimated OAS is the real OAS. If it's not the same, you have to try again. Exactly the same way as when you're solving for internal rate of return on a fixed cash flow bond. The OAS is the spread that, when added to the short-term rates in the model and used to discount back all the cash flows, gives you the market price.

In order to estimate duration and convexity, quite briefly, you take the current coupon Treasury yield curve and you essentially, not precisely, add a shock to that Treasury yield curve. You re-generate all the Treasury paths using the same correlation and volatility assumptions. You re-estimate all of the cash flows, which will now be different both because the coupon will be different and also because the prepayments will be different. Then you discount back those projected cash flows at the OAS you previously calculated. That gives you a new price. That gives you the price that would occur if the OAS stayed constant and the Treasury yield curve shifted by your shock. That's enough to get duration. If you want to get convexity, you have to do it again. You'll need to subtract a shock off the Treasury yield curve and go through the whole process. So, do you see what I mean? This is a very computer-intensive model, that's how we do it. The OAS model is a simulation, lots of rate paths, carefully constructed to make sure that it's an internally consistent model. You can get some weird answers if it's not. You have explicit volatility assumptions which are imperative for pricing both assets and liabilities. A model prices a complex security net of all the security's embedded options. It doesn't price embedded options separately. It prices the security net of all those embedded options.

So I've just outlined the theory and practice of our OAS models. The theory is elegant and self-consistent -- but so what! Does it work? -- that is the relevant question. Well, we didn't know for sure when we created this model in early 1986. We spent a long time testing it historically to make sure that it worked well. The tables are empirical test results and I'll describe how these tests were conducted. In this first test (Table 1), we examined the performance of the model over six-month holding periods, from the beginning of 1983 to the end of 1987. What we did is for each MBS that we were pricing, approximately 80 issues per six-month period, we calculated its option-adjusted spread and its option-adjusted duration using appropriate historical volatility levels. We then found what Treasury portfolio had the same duration; typically, one Treasury with a little shorter duration and one with a little longer duration. So we had the MBS and its duration-matched Treasury portfolio. We then calculated the historical total return of each MBS and its equivalent duration Treasury. We took the mortgage's total return minus the Treasury's total return. That's the differential return. We did not rebalance the portfolios over the six-month holding period. We would have had better results had we rebalanced those portfolios as needed during the holding periods. We took the top quarter in OAS across this whole five-year period and said, "Okay. What is our average outperformance and how often do those securities outperform? What was the percentage of outperformance?" Well, as you can see, if you took the top bucket, the top quartile of Government National Mortgage Association (GNMA) mortgages, and you made that your portfolio, you would have outperformed equivalent duration Treasuries by 258 basis points.

TABLE 1
Empirical Test Results
Historical Tests of Six-Month Holding Periods
1983-1987

Quartile	Average Outperformance	% Outperformance
	GNMA 30	
1	258 bp	97%
2	140	82
3	7	59
4	-123	28
	FNMA 30	
1	358	100
2	171	92
3	27	62
4	-31	46
	FHLMC 30	
1	387	100
2	203	92
3	9	65
4	-100	27

That's over five points a year! That's really remarkable outperformance. And 97% of the time you'll outperform; not quite a hundred, but you did very well.

What you're looking for in these tests is monotonicity. You want to see that the highest quartiles have the highest average outperformance and the highest percentage of the time when you outperform. And you see we get that. If you had actually bought the top quartile and shorted the lower quartile, you're getting almost four hundred basis points of outperformance every six months! As you see, we get extremely good performance out of these six-month periods. But it's not a great test because people were using poor valuation techniques in 1983, 1984, and 1985, so the market was very inefficient. It was easy to pick off value back then. So we did the test again on overlapping three-month holding periods, from the beginning of 1986 through the end of 1987 (Table 2). Now, these are raw three-month returns, and it's exactly the same study we did before, except now we've added a fourth column, the average change in spread. In other words, how much did the average option-adjusted spread change over that holding period? As you can see here, our first quartile of GNMAs, had an average outperformance of 135 basis points a quarter, 540 basis points a year. It outperformed 79% of the time. The market hadn't had as much time to correct value in three-month period, but that's still a high percentage. On average, those bonds tightened in option-adjusted spread by 25 basis points. In other words, they richened up. You bought them cheap and they got richer. That's what you want to have when you're long the security. We have perfect monotonicity in both the average outperformance, the percent of outperformance and the average change in spread.

TABLE 2
Empirical Test Results (Cont'd)
Historical Tests of Three-Month Holding Periods
1986-1987

Quartile	Average Outperformance	% Outperformance	Average Change in Spread
		GNMA 30	
1	135 bp	79%	-25 bp
2	94	75	-11
3	-18	47	15
4	-152	19	43
		FNMA 30	
1	150	86	-33
2	68	66	- 8
3	7	57	6
4	-83	24	42
		FHLMC 30	
1	158	89	-37
2	55	65	- 5
3	-11	55	12
4	-70	28	38

Well, let's take a little closer look now. Let's look at three-month holding periods (Table 3), but now with declines. So we're taking a much finer look. Well, here in the average outperformance we, again, get extremely good performance. The 161 basis points over the three months, well over six points per year. Down at the bottom we have underperformance of 140. We're still getting monotonicity, so even though you've divided these up into very small groups, you're still finding extremely good performance characteristics, and we're monotone in the average change in spread as well. Finally, we looked at one-month holding periods (Table 4). Now the market has only one month to correct, and here, finally, the monotonicity breaks down, but not by much. If you look down at the average outperformance, you get 49, 34, back up to 47, 17, 24. It's just remarkable that market value would correct itself over one-month periods. I didn't think we'd get results this good when we did this historical test. What these tests are saying is that the option-adjusted spread is an excellent indicator of value and that the duration estimates are unbiased. That means

TABLE 3
Empirical Test Results (Cont'd)
Decline Performance of OAS Model for
Three-Month Holding Periods 1986-1987

Decline	Average Outperformance	% Outperformance	Average Change in Spread
1	161 bp	92%	-40 bp
2	134	81	-27
3	110	73	-19
4	63	67	- 4
5	48	62	- 3
6	28	63	6
7	4	55	9
8	-63	39	25
9	-75	30	29
10	-140	11	57

TABLE 4
Empirical Test Results (Cont'd)
Decline Performance of OAS Model for
One-Month Holding Periods 1986-1987

Decline	Average Outperformance	% Outperformance	Average Change in Spread
1	49 bp	75%	-16 bp
2	34	68	- 6
3	47	73	-10
4	17	58	1
5	24	67	- 2
6	5	56	7
7	-1	55	4
8	-10	54	5
9	-20	53	9
10	-38	40	17

the duration might not be exactly right, but if we are a little bit long here, then we're a little bit short somewhere else. On average, we're getting the right duration. So, that's how option-adjusted spread modeling works on the asset side. That's how you can use it, and these are some historical results when applied to fixed rate mortgages.

MR. GRIFFIN: The purpose of my first example, a single premium deferred annuity (SPDA), is to further build your confidence with the option pricing techniques that Andrew has demonstrated. Most actuaries who are familiar with the SPDA product are aware that the policyholder has the ability to surrender the policy for a fixed amount regardless of where interest rates are. This represents a valuable interest rate option, and the cost of this option should be a component of the product pricing.

My second example will be a window GIC. I'll show how option pricing techniques can be used to value options that may be present in this product. I will also show how actuarial style analysis can provide a very useful breakdown of results.

Since interest rate volatility began in earnest back in the 1970s, there's been a lot of emphasis on measuring the market value of one's assets against the market value of one's liabilities. Measuring the market value of assets is relatively easy due to the secondary market for most assets. However, measuring the market value of liabilities is a lot tougher because there's no established secondary market for insurance products. In a large part this is the reason for the growth in popularity of the scenario testing approach.

Enter the required spread on assets, or RSA, a statistic that captures all of the traditional actuarial pricing elements as well as the value of any options inherent in the product. I will demonstrate how the RSA is calculated for my sample SPDA product. The RSA is the spread that must be added to the prevailing Treasury rates to discount the liability cash flows to the market value of assets. Initially, when we are in pricing mode, the market value of assets is the investable single premium remaining after up-front expenses. For interest-sensitive products, we'll need a set of interest rate paths like the ones Andrew has described to calculate this number.

Once we've calculated the RSA, we can calculate the interest rate duration of a liability. This is done by holding the spread constant and shocking the Treasury rate paths and then calculating new present values of the liability. The interest rate duration is the slope of the resulting present value curve. I will demonstrate later what an important statistic this is. Convexity gives us a second order estimate of present value sensitivity, where duration is the first order estimate. The mean term of liabilities is my own name for the classic Macaulay definition of present value weighted time to maturity, only averaged over the set of interest rate paths. Only for products where there are no interest-sensitive withdrawals and the credited rate on the product doesn't change will these two statistics be equal. In my example I'll try to emphasize how these two measures are different.

Before we get any further along with the option pricing approach, I'd like to make a quick comparison between option pricing and scenario testing. The two approaches are really fairly complementary. Both approaches start by generating a set of interest rate paths. In both cases we generate cash flows along each of our interest rate paths. With option pricing, only liability cash flows and not asset cash flows are projected. The liability cash flows will dictate the ideal characteristics of the asset portfolio. With option pricing, the set of interest rate paths must satisfy a number of arbitrage constraints. These constraints may seem like a real pain; however, when such a set of paths is constructed, the short-term rates can be used to discount cash flows back to their market value. I hope my presentation will demonstrate how valuable such a set of paths can be.

To understand how this approach works, let's consider first a couple of very simple cases (Table 5). First let's imagine that our insurance company can issue current on-the-run Treasury bonds at their market price. The first line shows the issue of a 1-year Treasury bond. The first step is to generate a set of Treasury interest rate paths. Projecting the cash flows from the bond is easy because they're fixed. In this case, if we add a spread of 0 to the short-term Treasury rates, the cash flows discount to the market price, which is $100; hence the RSA is zero. The interest rate duration and mean term of this liability are the same, equal to that of a one-year bond.

TABLE 5

Pricing Steps

	Required Spread on Assets	Interest Rate Duration	Mean Term
1 Year Treasury	0 bp	1.0 yr	1.0 yr
3 Year Treasury	0	2.6	2.6
Bare-Bones SPDA - No expenses - No commissions - No surrender charges	0	2.6	9.7

If, instead, we assume the issue of a three-year Treasury, and go through the same calculations, the RSA is again zero but the interest rate duration, and mean term, increase in accordance with the three-year bond. The RSA of 0 makes sense because if we assume no expenses and no profit target to match this liability, we would not require a spread over Treasuries.

Further imagine that we will now issue what I've called a bare-bones SPDA. We still have no expenses, no profit target, and no surrender charges. We will guarantee the current three-year Treasury for three years, and at the end of three years we will reset the rate to the then prevailing three-year Treasury rate, for the next three years and so on. There is a base lapse rate, and gradually the business will run off the books. Modelling this product and discounting the cash flows gives an RSA which is still zero. This makes sense because we're still crediting Treasury rates. The interest rate duration is the same as the three-year Treasury bond, which would seem to make sense based on our credited rate strategy and rate reset timing. There is now a much longer mean term of liabilities because the business is "on the books" for a much longer period.

The first line on Table 6 is a repeat of the last line on Table 5. The next line is the first of a number of pricing steps I'm going to go through. First I'll introduce some "up front" expenses. The new RSA tells me that I'll have to earn 87 basis points above Treasuries to recoup these expenses. In fact, even though this number appears to come out of a black box, the shortcut to estimate this number is very intuitive. One would express the upfront expenses as a percent, and in effect amortize them by dividing the percent by the mean term of liabilities, to get 87 basis points.

TABLE 6

Pricing Steps

	Required Spread on Assets	Interest Duration	Mean Term
Bare-Bones SPDA	0 bp	2.6 yr	9.7 yr
Deduct $2,300 of Up-Front Expenses from Market Value of Assets	87	2.6	9.4
Add Annual Renewal Expense of $100	124	2.6	9.4
Credit 75 bp Below 3-Year Treasuries	49	2.7	9.4
Collection of Surrender Charges	23	2.7	9.4
Include Interest Sensitive Withdrawals	51	2.1	7.5

On the next line I've introduced an annual renewal expense. As we would expect, the RSA increases by the annual expense, expressed as a percent. For both of these expense additions, there is no appreciable change in the interest rate duration target.

On the next line, instead of crediting and recrediting three-year Treasury rates, we'll assume that we'll always credit 75 basis points less. Accordingly, the RSA decreases by 75 basis points to 49 basis points.

Another aspect that will make our product "cheaper" is that we'll collect surrender charges when policies are surrendered. Once again, this has a cost effect, but no appreciable duration effect.

This next step is the most interesting. Here we introduce interest-sensitive lapses. I've assumed that whenever our credited rate falls too far below prevailing new money rates for SPDAs, there will be interest-sensitive withdrawals on top of the base lapses. This calculation takes into account the policyholder's presumed amortization of any surrender charges. I hope these tables show that the RSA is a sort of "all-in-cost" measure and provides one consistent approach to the whole exercise.

You can see from the increase in the RSA for interest-sensitive withdrawals that there is a "cost" to this phenomenon. This also produces a lower interest rate duration, for the same reason that adding a call provision to a noncallable bond decreases its interest rate duration. Also, because we have faster lapses now in aggregate, the mean term of the liabilities is also decreased.

Up to this point, we've assumed a total reset to the new prevailing three-year Treasury, every three years. Suppose that instead (Table 7) we were to reset our credited rate half of the way between the previous rate and the new prevailing three-year Treasury. Those results are shown on the second line. The increased RSA shows that this is a more expensive strategy to follow in this case. Because we reset only a part of the way to new three-year rates, the interest rate duration has lengthened. The product now behaves like a longer-term fixed rate instrument than before. Because the product's credited rate does not track new money rates as closely as before, there will be more interest-sensitive withdrawals and the mean term of liabilities is decreased.

TABLE 7
Rate Reset Strategy

	Required Spread on Assets	Interest Rate Duration	Mean Term
100% Reset	51 bp	2.1 yr	7.5 yr
50% Reset	68	3.0	6.9
Downward Bias - 50% up -100% down	43	2.2	6.5

The bottom line shows the effect of following what I chose to call a "downward basis" strategy. Another possible description would be "bait and switch." Under this approach, we would reset 100% of the way to new money rates if that adjustment is down, but only 50% of the way to new money rates if that adjustment is up. As one would expect the downward bias strategy is cheaper. I'm not necessarily a proponent of this particular strategy, I just want you to be aware that these more dynamic strategies can be modelled easily in this framework. The relative cost of these strategies will depend on the interest-sensitive withdrawal algorithm used.

What does the RSA actually get used for? The RSA is the first element in calculating our spread target for assets. Once we have the RSA, the next step is to add any asset-based expenses. These typically include at least investment expenses and might also include a charge for the default risk of the assets. The final element is, of course, profit. This is such a touchy subject I didn't even want to make up a number for this, so I just used the algebra technique of representing it by a letter.

The second target, the interest rate duration target, is the one that should be used for purposes of asset portfolio management. This is a very important statistic, as I'll show soon. What these statistics tell us is that our liability will behave like an asset with these characteristics. If we wanted to further define the liability, the next step would be to calculate a convexity measure.

Let's go back to the scenario testing approach, where we model an investment strategy at the same time as modelling the liability. For each investment strategy we get a surplus result for each scenario, and we can compare them using a bar chart like the one shown (Graph 14). We chose a spectrum of rather simple investment strategies ranging from short term (second from the left) to long term (on the right). With the scenario testing approach we would typically keep testing different strategies until (a) we got exhausted, (b) we used up our budget of central processing unit (CPU) time, (c) we happened to find a strategy that met our constraints.

GRAPH 14

SPDA Investment Strategies

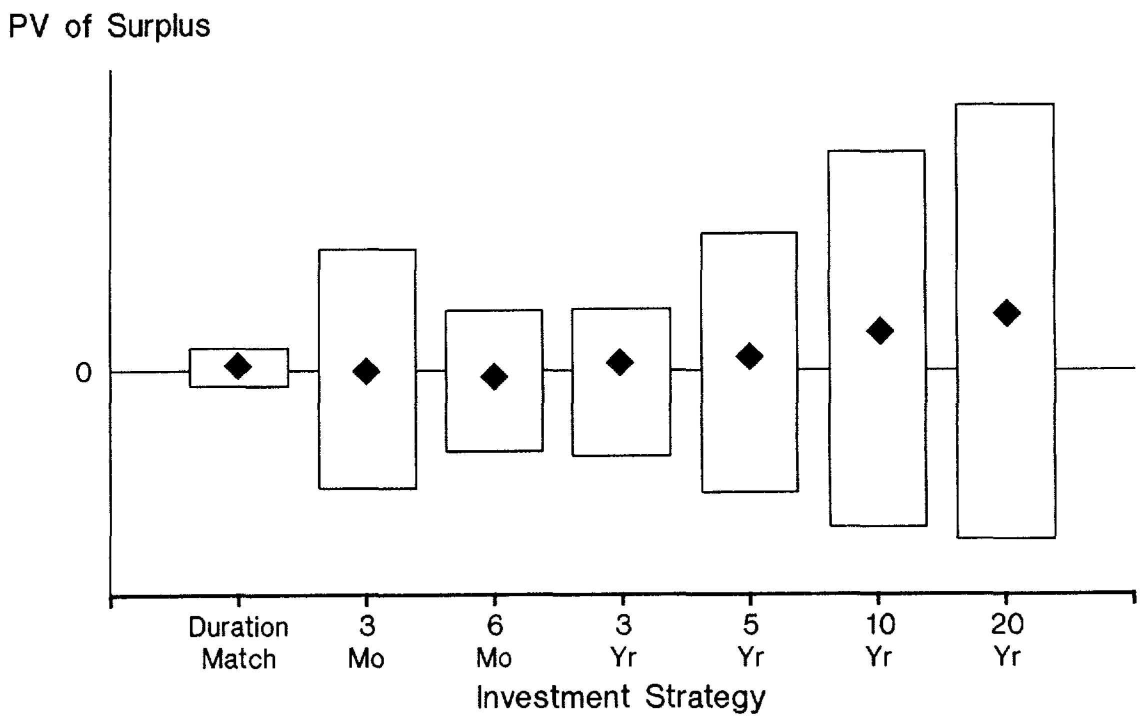

In contrast, the option pricing approach is not a trial and error approach. The characteristics of the liability are determined, and those become the targets for the asset portfolio management.

To check how powerful this technique is, we did a glorified scenario testing run, where the investment strategy was to maintain a proper interest rate duration match between assets and liabilities. The results of that run are shown by the small bar on the far left of the graph. I think it is safe to say, that if you are risk adverse, the duration matching results are very good results. These results were not easy to come by. Along each path, at each point in time, an interest rate duration for the liability had to be calculated. This entailed generating a set of interest rate paths, shocking them up and down, projecting future cash flows, and so on. In total it involved the generation and evaluation of more than one million interest rate paths! Recall from Bob's talk that matching durations is a very good first step, but it does not provide one with iron-clad protection. My example incorporates duration matching only, not convexity matching. I'm sure the bar would be even smaller if we were able to add convexity matching, but our resources were limited. My next and final example will be a window GIC. I'm not going to drag you through all the pricing steps again, I'm just going to focus on the option-like characteristics this product can have. My sample window GIC has a considerable amount of what I call cash flow antiselection and is probably representative of the riskier end of the spectrum of GIC cases. My example is meant to show how these option characteristics can be evaluated and what can be done after the fact to analyze experience.

The first aspect of this antiselection is deposit antiselection (Table 8). Suppose that we have a window GIC where we expect to receive a million dollars per month throughout the calendar year. However due to certain plan provisions, or lack thereof, we know that we should expect to receive less money if rates go up, and more money if rates go down. I have expressed this tendency in the deposit grid. How plan provisions affect cash flow antiselection comes under the heading of GIC underwriting, which is a discipline in itself, and I don't have time to even really touch on it.

TABLE 8
Expected Deposits

Change In Interest Rates	January	February	....	November	December
+1.0%	$ .85	$ .85	....	$ .85	$ .85
+ .5	.92	.92	....	.92	.92
0	1.00	1.00	....	1.00	1.00
- .5	1.07	1.07	....	1.07	1.07
-1.0	1.15	1.15	....	1.15	1.15

The next aspect of antiselection is withdrawal antiselection (Table 9). The GIC pays withdrawals at book value. We have an expected withdrawal level of 10% per year; however, we also expect antiselection in the form of higher withdrawals when rates have gone up and lower withdrawals when rates have gone down. This tendency towards antiselection will result from certain plan provisions and/or lack of other plan provisions. The last two Tables are examples of assumptions of the partial exercise of options. For example, an assumption of a totally economic exercise of the book value withdrawal feature would be that withdrawals are zero if rates go down and go quickly to 100% if rates go up.

TABLE 9
Expected Withdrawals
(Annualized)

Prevailing Interest Rates Minus Blended Rate of Plan	Year 1	Year 2	...	Year 5
+2.0%	12.5%	12.5%	...	12.5%
+1.0	11.3	11.3	...	11.3
0	10.0	10.0	...	10.0
-1.0	8.7	8.7	...	8.7
-2.0	7.5	7.5	...	7.5

Table 10 shows the risk charges calculated for these 2 types of antiselection. These charges are determined by projecting deposits and withdrawals over our set of many interest rate paths. The results are very dependent on the assumed volatility of interest rates that goes into constructing the set of paths. It's also possible to incorporate into the model the effect of various contract provisions such as caps on the deposit amount.

TABLE 10

Antiselection Risk Charges

	Risk Charge
Deposit Antiselection Risk	.18%
Withdrawal Antiselection Risk	.13
Total	.31%

Now I'm going to shift gears a little bit with my same window GIC example and consider the situation at the end of the window period. On a case like this where we expect a certain amount of antiselection we could try to set up a hedge by using our risk charges to buy put and call options. It probably wouldn't be possible to construct a perfect hedge, but we could put something together.

I'm assuming here that no particular antiselection hedging was done. This is very common, and I like to call this strategy "pocket and pray." The insurer pockets the risk charges and prays that low interest rate volatility and/or conservative underwriting will turn part of the risk premium into profit.

Table 11 shows a month-by-month breakdown of the deposit experience under the window GIC. I've assumed that the insurance company took a loan to raise the money to buy the bond to back the GIC. Repayment of the loan is scheduled as one million per month through the window period, which is identical to the expected deposits if rates don't change. So for each month, Table 11 shows where interest rates were that month and what the actual deposits were. The net cash flow is then calculated and the profit or loss determined based on whether bonds had to be bought or sold and at what yield level.

TABLE 11

Deposit Antiselection Experience

Month	Change In Interest Rates	Actual Deposits	Loan Repayment	Net Cash flow	Profit or Loss
January	+1.0%	$ .80	$ 1.00	$-.20	$-.008
February	+1.5	.90	1.00	-.10	-.006
March	+ .5	1.15	1.00	+.15	+.003
.	.	.	.	.	.
.	.	.	.	.	.
.	.	.	.	.	.
.	.	.	.	.	.

In this example, the total loss from deposit antiselection, when amortized over the life of the case, came to six basis points. Since our pricing assumption for this risk was 18 basis points, simple math tells us that we made an additional profit here of 12 basis points. That seems easy enough, and maybe my analysis should end here, but it doesn't.

I'd like to decompose that 12 basis points of profit into 2 separate sources.

The first step to do this is to repeat my calculation of the actual loss incurred (Table 12), except instead of using the actual deposits received, use what we would expect to receive, given where interest rates were that month, in accordance with the grid we prepared of expected deposits. This allows us to calculate our expected deposit antiselection loss, knowing what happened to interest rates.

TABLE 12

Expected Deposit Antiselection
Knowing Interest Rate Path

Month	Change In Interest Rates	Expected Deposits	Loan Repayment	Net Cash Flow	Profit or Loss
January	+1.0%	$ .85	$1.00	$-.15	$-.006
February	+1.5	.78	1.00	-.22	-.013
March	+ .5	.92	1.00	+.08	+.002
•	•	•	•	•	•
•	•	•	•	•	•
•	•	•	•	•	•
•	•	•	•	•	•

Fourteen basis points is the result of calculating the expected deposit antiselection loss, knowing the path that interest rates actually took. Subtracting the actual loss from this gives us the profit from misestimating antiselection. In effect, knowing where rates went, if GIC depositors had behaved the way we originally predicted, our loss would have been 14 basis points and not 6 basis points.

The other component of the 12 basis points is the effect of interest rates having been less volatile than we projected when we did the pricing. We calculate this component by subtracting the 14 basis points from the pricing assumption.

I hope you'll agree that this can be a very useful breakdown. For example, it's quite possible that one of these components could be a profit and the other could be a loss. To get a better feel for this, imagine the case where we don't expect any antiselection. Any profit or loss in this case will fall under the category of antiselection misestimation. In a sense the profit or loss from antiselection misestimation is like our underwriting score card. It measures whether we over or under estimated deposit antiselection. The profit or loss from "volatility misestimation" measures, in hindsight, whether or not it was a good decision to follow the pocket and pray strategy. If it were possible to construct the perfect hedge, and we followed that strategy, this component would be zero.

In summary I want to make two points. My first point is certainly not a new one. When insurance products have valuable options in them, they should be priced. You don't want to give these things away!

Point number two is that actuaries must study the utilization of interest rate options by policyholders. The corporate Treasurer's exercise of call options in his or her company's bonds, can and is, presumed to be totally economically efficient. In that situation the decision is made by one very informed and financially aware individual who is potentially assisted by many other sophisticated financial people. However, in the situation of MBSs, one must model a less than totally efficient exercise of the prepayment option. There are still plenty of mortgages outstanding that are at higher than current rates.

For insurance products, the task is very similar to that of mortgage backed. We have to estimate the degree to which policyholders will exercise their options. I think most actuaries would concur that an important variable for individual products could well be the distribution system used to sell the product. And, as important as making estimations, we as actuaries must remember to look backwards and calculate how accurate our projections were. Another way to think about this is that the utilization of policy options is now an important assumption. We have to make sure that we remember to do an actual versus expected calculation here, too, just the way we were taught to do one for mortality experience.

MR. STEVEN P. MILLER: I have a question for Mark. Have you done any studies on the correlation between maximizing option-adjusted spread on your assets minus the required spread on your liabilities and the present value of dividendable cash flows? I mean dividendable as in paying dividends to your stockholders, since that's eventually what you want to do, to maximize the amount of dividends that you can pay to your stockholders. That, unfortunately, requires a relationship between the investment income on your assets and your taxes and when you realize your profits and things like that.

MR. GRIFFIN: Even though my presentation made no reference to statutory accounting and taxes, it doesn't mean that I've totally forgotten them. Our typical assignment is to look at an existing portfolio of assets in terms of how it might be restructured or optimized. There is typically a duration constraint as dictated by the duration of the liabilities. There are typically minimum and average credit constraints on the "new portfolio." There may also be constraints as to capital gains or losses that may be incurred and to book income. So we want to widen the option-adjusted spread, but subject to a number of constraints. Even though what I've shown here appears to be a purely economic exercise, it is possible to do this type of study and incorporate statutory accounting and taxes.