2002 Valuation Actuary Symposium
September 19–20, 2002
Lake Buena Vista, Florida

Session 6TS
Asset Modeling Concepts

Instructor: Scott D. Houghton

Summary: The instructor explains how to scrutinize model results and use models to make strategic decisions. Specific topics include validation of asset models; modeling considerations for derivatives and assets with embedded options; using models to assist with portfolio management; and using models to improve portfolio yields and realized product spreads by taking efficient risks without betting the company.

MR. SCOTT D. HOUGHTON: I'm vice president and consulting actuary at AON Consulting in Avon, Connecticut. Before AON, I worked in asset/liability risk management for North American Security Life, which is now part of Manulife. Previous to North American, I did product development and valuation work for Boston Mutual Life Insurance.

The session is divided into four parts. There's a brief introduction. I'm going to talk about validation of asset models and then about modeling considerations for assets that contain embedded options. I will conclude with methods for using models to assist with portfolio management.

One of the key things we try to do in our models is strike a balance between a model that oversimplifies things and one that tries to incorporate too much and isn't practical to use. We want to avoid oversimplification, but we also want to avoid too much complexity. Since the symposium is at Disney World, let's take the story of Pocahontas, the Native American who

Copyright © 2003, Society of Actuaries
Note: The chart(s) referred to in the text can be found at the end of the manuscript.

helped broker peace between her tribe and the English colonists at Jamestown. The real story has a pretty complex love triangle between the major characters or maybe even a love rectangle depending on what version of history is correct. The Disney cartoon movie version leaves most of that out. There's a good reason for the oversimplification. It's a lot easier to sell Pocahontas coffee mugs in a gift shop that way.

Another example is the Alamo, the building in San Antonio, Texas that was originally built as a mission but used by the Republic of Texas' army in its battle for independence with Mexico. You know Ron Howard—the director of last year's best picture, Oscar winning movie, *A Beautiful Mind.* The younger actuaries in the room think of Ron Howard as a director. The older ones think of him as Richie Cunningham on "Happy Days." People like myself think of him as Opie from "The Andy Griffith Show." Ron Howard has been hired by Disney to direct the movie version of the Alamo, and he has had some differences of opinion with Disney over the casting and the script for the movie. Ron Howard did get some criticism for leaving out a few things in *A Beautiful Mind,* but now his story is that he's on the side of history.

In the published version of history, the villain in the story is General Antonio Lopez de Santa Anna who commands the Mexican army. The good guys consist of the Republic of Texas army who is fighting for its independence from Mexico. Most of you know the story and the battle of The Alamo. Davy Crockett and 188 others are killed when the Mexican army refuses to accept their surrender. The good guys lost, but their treatment by the Mexican army angered and motivated the rest of the rebellion in Texas.

General Santa Anna had a reputation for being a big micro-manager and made virtually all of the decisions for his army. Then he got distracted. There were two big distractions. He developed a drug habit and became addicted to opium. The second distraction was a young woman named Emily Morgan. Ms. Morgan was an indentured servant who became a prisoner in one of Santa Anna's campaigns against the Texas army. So now Santa Anna's Mexican army, who is used to being micromanaged, were without leadership because their leader was either in Ms. Morgan's company or under the influence of drugs most of the time. Sam Houston, who commanded the Republic of Texas army, realized he had an opportunity there, and he launched a successful surprise attack and captured General Santa Anna at San Jacinto.

Since I should probably start talking about asset modeling soon, I'll try to get to the point here. Texas gained its independence and eventually became a U.S. state. We know who Davy Crockett is. Sam Houston and Stephen Austin, the heroes of the Texas army, have Texas cities named in their honor. General Santa Anna faded into historical obscurity, and Emily Morgan, for her role in Texas's victory, was released from her indenture and moved to New York. Historians lost track of her from there, but she was immortalized in a folk song written about her called "The Yellow Rose of Texas."

So, with the Alamo, we have a real story that's full of violence, drugs, and sex. The real story is probably rated R. The good guys lose, at least at the Alamo. That doesn't really sound like a typical Disney film, but the PG-13 version from Disney Touchstone Pictures will premier in July of 2003. If there's any of you out there with PG-13-rated asset models, we'll talk about what we can do to add some more analytics and some more sophistication to them without adding a great deal of work or complexity.

Let's discuss the main things to consider in validation of our asset models. The first one is pretty obvious. The threshold of our validation depends upon the purpose of our model. This is the case where we need to strike a balance between how much time, effort, and resources we spend versus the results that we get. Validation is an exercise of what an economist would call decreasing marginal utility. The first block of time we spend on validation adds value to our model and our results. The next block of time adds more value, but it adds less value than the first block of time did. The more strategic the decisions are that we make with our models, the more comfortable we would like to be with what comes out of them, and the more time and effort we can justify spending on validation.

If the purpose of our model is cash-flow testing to satisfy a regulatory requirement, then we still want to validate the model to confirm that the input and output of our model is reasonable.

However, if we're not using the results to make strategic decisions, then the value of a marginal or the additional hours spent on validation is lower.

There are two broad categories of validation. The first type is what I'd call static validation. Static validation involves examination of balances and starting amounts. We make sure that the balances that are in the model at the valuation date are in sync with our company's portfolios on the valuation date. The second type of validation is what I call dynamic validation. Dynamic validation involves comparison of the projected data that comes out of our models to actual financial data.

Once again, the first type of validation, static validation, involves examination of balances and starting amounts. We make sure the balances that are in the model at the valuation date are in sync with our company's portfolios on the valuation date. It's important to validate both input and non-input items. We'd input book value and PAR value for bonds. The input and output of our model should match. An example of a non-input item we can validate with is market value. If our modeling software projects future cash flows for a bond and discounts them to determine market value, we can check what the model produces for market value against a published source for the market value on the valuation date. If these two market values don't line up, something in our model isn't quite reasonable. There's either a problem with the projections, the cash flows being projected, or with the discounting or possibly something else.

Sometimes the modeling software takes market value as an input item but might project discount spread over Treasuries needed to get that starting market value at the valuation date. So, in that case, we can compare the discount spread that the model projects to what's available in that marketplace for that type of bond and see if that's reasonable. Another example of a non-input item we can use for validation is book yield. If we get the book yield of a bond correct, then it's very likely that all of the parameters that determine the cash flow and the statutory values for that bond are correct. Static validation can be done at an individual security or at a Committee for Uniform Securities Identification Products (CUSIP) level by asset class and at a portfolio level.

Dynamic validation involves comparison of the projected data that comes out of our model to actual financial data. Dynamic validation involves an elapsed time period. We can look at the actual cash flows and investment income that come out of the model over a period and compare this to the actual cash flow and investment income that come out of the portfolio for the same period. To effectively do dynamic validation, we need a model that was created in the past. If we happen to have a model that was created in the past, then we're in good shape, but if we don't, we can create a model with current data and then wait six months or so. We'll have a model that was created in the past. We can dynamically validate prospectively or retrospectively.

The problem or the challenge with dynamic validation is that it's a lot of work. If you don't want to devote the resources necessary, we can use a simplified or a shortcut version. I'm going to talk about some simplified techniques that can be used for dynamic validation that really involve taking a top-down approach. We can start with our company's annual statement and a model that was created as near as possible to the date of the annual statement. In the formula below, I is investment income. A and B are the start-of-the-year and end-of-the-year book values. For a company's entire asset portfolio, we can calculate a book yield for that portfolio for the year 2002 using our 2002 annual statement data that will be available in four months or so. If you have last year's cash-flow testing model, you can do the same calculation for 2002 and compare the answers.

$$2I / (A + B - I) \qquad \text{[Equation 1]}$$

There are a couple things to consider when we do this. Chances are the actual interest rate scenario that has occurred in 2002 probably doesn't match any of the scenarios you used to run last year's cash-flow testing model. Interest rates were pretty level January through June. So, if we did want to take more of a shortcut and validate a six-month period, January through June, we could look at last year's cash-flow testing model with a level scenario. To look at a full year or a longer period, we'd need to consider that long-term Treasuries went up about 35 basis points or so in March, dropped back to their start-of-the-year level in May and then dropped 1% in June, and additional 1% in July, while short-term Treasuries have remained pretty much the same all year.

If you use a level scenario for cash-flow testing validation for a period longer than six months, then we'd need to consider the yield curve and make an adjustment for assets purchased in July or later. A second adjustment that would typically be made is for new premium. Cash-flow testing models don't typically include new business. The annual statement data will reflect the new business that has come in, and those funds have probably earned a lower rate than the rest of the portfolio, given the current interest rate scenarios. A third adjustment would be in a case where a company has experience that's materially different than the assumptions that were used in the model. A fourth adjustment would be if a company practice with respect to investments and investment policy is different than the assumptions in the model. Fifth, there might be some more specific adjustments that would apply to particular companies in specific situations. But I don't think these adjustments are too burdensome.

So I think it's a pretty straightforward dynamic validation exercise to validate last year's cash-flow testing model to actual 2002 data. Once that's done at a company level, the exercise can be repeated at a portfolio level, and by portfolio level, in this case, I'm talking about very broad asset categories that you'd see in annual statement Schedules B and D that hold information for mortgages and bonds. Again, it would still be necessary to make those five potential adjustments that I talked about. We can also use price curves as a validation tool for assets with embedded options, and I'll talk a little bit more about that later on.

We've talked about using annual statement data. This exercise can also be done with investment accounting system data at a lower, finer level. We might want to do that anyhow, but we would certainly want to do that if there are any discrepancies or unexplained differences after we did it with the annual statement data. Typically, if there are unexplained differences, we'd want to narrow those down to an asset class and eventually an individual security level. When we do look at individual securities we can, of course, use book yields and market values or the discounts over Treasuries implied by those market values as a reasonableness check. A technique we could use to make this easier is to just take electronic output from our models. So we run our model, get output that has a CUSIP in one column and book yield in the other column, and then we can get annual statement Schedules B and D electronically from our accounting folks so we have a second set of data that has the CUSIP and the book yield in the two columns. We paste both of the schedules into a spreadsheet and sort them by CUSIP and just take a look at the assets that have material differences.

A couple of considerations if we do that are that the yield for mortgage-backed securities and collateralized mortgage obligations (CMOs) will depend on the prepayment assumption that's in our model. If these securities are bought at a discount from PAR, then if we use a higher prepayment assumption, the discounts will amortize faster, and the yield will be higher. The converse is true if these securities are bought at a premium from PAR. So we want to have an idea of what the prepayment assumption was for the yield and the data that we were validating to. Sometimes for Blue Book purposes, the securities are amortized over their stated maturity, which means, effectively, that there are zero prepayments. So if a CMO is purchased at a discount from PAR, then we would get a higher yield in our model than the Blue Book in that circumstance. A second consideration is that I've assumed that the assets are modeled on a seriatim basis. If they're modeled on a portfolio or a grouped basis, then we need to do this type of validation at a higher level.

An additional tool that's available for validation is a price curve. It can also be called a price behavior curve. Many of you are probably familiar with using price curves as a risk management tool. I'm going to talk about using the same tool but for a different purpose — as a validation tool. The way we create a price curve is to plot a series of parallel interest rate shifts on the horizontal or the x-axis. On the y-axis, for our independent variable, we're plotting the market value of our asset or asset portfolio. We can also plot a separate line for liabilities, but I'm just going to be considering assets.

Chart 1 is an example of a price behavior curve for a mortgage-backed security. The horizontal or x-axis shows a parallel interest rate shift. A shift of zero means that just the interest rate scenario is level. One percent means that we do an immediate shift or immediate increase in interest rates of 1% and then leave them level for that. Two percent means an immediate shift to 2%. Negative 1% means an immediate drop of 1%, and so on. Many of you are probably used to seeing zero in the center of the graph, but the current yield curve is very low, especially on the short-term end. To effectively plot this, I need to put zero on the x-axis to the left of the center. We have the interest rate shift on the x or the horizontal axis. We have the market value on the vertical or the y-axis. We might do this exercise for a portfolio of assets and liabilities and look at the market value under each of the interest rate shifts and use this as an interest rate risk management tool. The big criticism or the weakness in this, as an interest rate risk management tool, is simply that it only measures risk for parallel shifts in the yield curve.

The yield curve movements generally don't consist of parallel shifts. It doesn't really help us if the yield curve changes slope or changes shape. So we can use it only to manage duration risk but not convexity risk. I'm just using this tool for different purposes. I'm using it for validation.

Chart 1 is a price curve for a mortgage-backed security (MBS). In this example, this curve tells me that my prepayment assumption is acting somewhat reasonably. It tells me the model is really working the way I intended. I know that because of the shape of the line. The bend in the line or the negative second derivative or negative convexity is typical for any type of bond that has a call or a prepayment option. If interest rates increase, there are two bad things that happen to this bond. First, interest rates have increased, so the market value goes down because we're using a higher discount rate to value the market value of the bond. Second, the term of my mortgage-backed security extends when interest rates increase because fewer people will prepay their mortgage loans.

Within the range of the middle of the page here, with each marginal increase in interest rates, the market value decreases a little bit more than it did with the previous marginal increase in interest rates. We see a little bit of a bend in this curve. A corporate bond or Treasury bond without any prepayment or call provisions would be very close to a straight line. Since we do have a prepayment assumption, it works to our disadvantage. The market value drops faster when rates increase, but it increases slower when rates decrease.

The most common sources of modeling data that I run into include bad data within the investment accounting system that becomes our source data for our asset models. An example of assets that are misunderstood or misclassified would be an asset-backed security that's modeled as a bullet bond. That would also include assets that don't always fit neatly into one of the modeling classes used by our software. If our software expects accrued interest and it's missing, then our software might overstate the yield and the investment income for a bond. Timing differences are very important or sometimes can be very important.

For example, say we have a bond where the coupons are payable semi-annually, on June 30 and December 30, but the stated maturity of the bond isn't June 30. Say June 30, 2010 falls on a Sunday. The stated maturity of the bond is July 1, 2010. If we tell our software that maturity is July 1, and it incorrectly models coupons in January and July instead of June and December, and we do a year-end model valuation with December 31 data, then our software thinks coupons are payable in January and July. It would expect that the December 31 value should really have, between five and six months accrued interest in it, but in reality, since our coupon was payable in December, our December 31 coupon has already been paid, so the actual accrued interest is close to zero.

If we model a bond like this, our book yield will get distorted. The solution is pretty simple. That's just to tell our modeling software that the bond matures in June of 2010 instead of July 1, 2010. If we're electronically getting some data from an investment accounting system, we would want to have the procedural infrastructure in place to really correct that problem or that challenge for us permanently. That would be a lot more preferable than trying to correct it each time we do a modeling cycle for the next eight years or so. That concludes the portion of the talk on validation.

The next section deals with modeling considerations for derivatives and assets with embedded options.

The main concern here is capturing the key risks of the assets within our model. We want to make sure that our model reflects the credit risk, interest rate risk, equity risk, prepayment risk, and volatility risk that our assets might exhibit.

There are two main types of models that are used to model call and prepayment risk. There are interest rate or spread-based models, and there are market value or price-based models. Spread-based prepayment models compare the interest rate being paid on a bond with an interest rate that's available on a newly issued bond. It would call the bond if the new interest rate is sufficiently lower. Rates do need to be adjusted to reflect call premium. If our model doesn't do that, we'll have a little bit of danger in having a PG-13 rated asset model in that case. Price-based models compare two market values: the market value of the bond, assuming it's not called, and the call price of the bond.

Sometimes a bond and our mortgage prepayments are simplified and only depend on one variable, which is interest rates. Sometimes, the borrower's behavior depends on more than just interest rates. To have a bond called or to have a mortgage prepayment or a mortgage-backed security exhibit prepayments, the underlying borrower has to have both the ability to refinance and the incentive to refinance. To have the ability to refinance, they need to be creditworthy. If it's a corporation with callable debt, the rate they can refinance at depends on the credit-worthiness, or their ability to refinance and save money depends on their creditworthiness. If we're looking at a commercial mortgage, we need to look at the current loan-to-value ratio. If the property has lost value due to vacancy or general market conditions, the borrower doesn't have the ability to refinance.

If we're looking at a mortgage-backed security or a CMO where the underlying collateral is a residential mortgage, the borrower might not be able to refinance favorably because their credit is not as good as when they took out the loan. It can even be fairly minor credit issues that make things harder for people to refinance. If a borrower has taken out a second mortgage, it can make things a challenge in certain situations. There are some second mortgages that are pretty restrictive, and they become first mortgages if the original first mortgage was prepaid. The borrowers are really looking for a second mortgage to replace their first and second mortgages, and the terms in second mortgages aren't as favorable. Or, the borrowers are looking for a high loan-to-value ratio first mortgage to replace their existing two mortgages, again with terms that might not be as favorable as a conventional or a conforming fixed mortgage.

In today's market, many of the housing markets in the U.S. have prices that have increased. Borrowers have tapped into this equity a little bit. They might have a conforming first mortgage loan, but they would now need a jumbo loan to replace it, and to refinance it. Those loans are at a quarter or three-eighths of a percentage point higher than a conforming loan. That reduces the borrower's incentive to refinance. We can use prepayment as a tool in validation of our models. I'm in the second part of the talk, and I'm supposed to be talking about modeling assets with embedded options. I'm going to take a jump backwards and point out that if there are assets that our models say should be called or prepaid, and in reality they're not, then we've quite possibly missed something within our model.

Liquidity is also a big issue as far as incentive to refinance. On corporate bond debt, if a company wants to refinance the debt, there's a cost to doing this. There is a cost related to calling the old debt, and there's a much higher cost that's paid to investment bankers to issue the new debt. Commercial mortgages and residential mortgages are similar. Although the costs for residential mortgages tend to be declining over time, the market is becoming more liquid. There's legislation in front of Congress now that might make this even more liquid by making the charges that go with residential mortgages more transparent.

In summary, bond calls depend on both the ability and the incentive to refinance. The rates at which bonds are called depend on the credit rating of the issuer, on the level of interest rates, both the Treasury curve and the relevant corporate bond curve for that issuer, and, of course, the underlying economic environment. It's fairly common in our modeling to simplify this complex prepayment assumption to a simpler one that depends only on interest rates or the prices of the bond, you know, with and without the call option considered. These two models I talked about earlier are the spread-based and the price-based prepayment models.

Mortgage prepayments, of course, depend largely on these same issues. There are economic and noneconomic prepayments. Bonds, on the other hand, are just motivated by economics. There are seasonal prepayments, and they're motivated both by economics and investor behavior. I really mean borrower behavior and not investor behavior. Investor behavior can have some influence on prepayments. Mortgage servicers can be either the lender or an agent of the lender that collects and processes the payment and services the mortgage. It's difficult to establish a direct statistical link, but there are some prepayment studies that suggest that certain mortgage servicers have higher prepayment rates than others. The main driver is that their policies differ with regard to grace periods and late payments. Some servicers engage in telemarketing and annoy the borrowers and give them an additional incentive to refinance.

One final piece is that acquisition costs for mortgage lenders to acquire a mortgage are fairly high, and there are a lot of lenders that aren't real happy with the trend that mortgages are becoming more and more liquid. The terms of the mortgages get shorter, but the high acquisition costs stay the same. Because of these two factors, residential mortgages that have prepayment penalties are becoming increasingly common. Companies that invest in mortgage-backed securities or collateralized mortgage obligations (CMOs) are being offered deals that include these types of mortgages in the mortgage pool. Of course, we're asked to reflect this in our modeling. Most of these prepayment penalties for residential mortgages are simplified. For example, a borrower in a commercial mortgage might have a make-whole provision where the prepayment penalty is a function of interest rates. Typically, a residential borrower might just have a prepayment penalty that's a lot simpler to understand, and that's either a fixed percentage of the mortgage loan balance or a declining percentage of the mortgage loan balance.

The key words there are simpler to understand because it's a lot simpler for the borrower to understand. It's a lot more difficult for us to model. Those prepayment penalties don't apply to small amounts of additional principal that the borrower might send in every month. They also typically don't apply if the borrower sells the property. They only apply to refinancing. In these types of prepayment penalties, the modeling can get complicated. Borrowers could move to a different house to effectively refinance their mortgage, although this makes the cost or the threshold for them a little bit higher. Therefore, the liquidity might be a little bit lower. You could also have two neighbors that sell each other their houses and lease them back from each other. So, they can effectively refinance that way.

I'd like to discuss an example of investor behavior with regard to mortgage prepayments. I'm referring to residential mortgages without any prepayment penalties here. There's two interest rate scenarios. There's Scenario A that declines linearly from Periods 1 through 3, and then it declines from Period 3 to 4 a little bit less than it did in the previous periods. There is also Interest Rate Scenario B that declines for three periods, same as Scenario A, but then it increases slightly from Period 3 to Period 4. If I have a mortgage that's taken out at time equals one, and interest rates drop through Periods 2 and 3, of course, I'm going to see rising prepayments in Periods 2 and 3.

What happens in Period 4 with most residential or mortgage-backed securities or CMOs is that in the case of something like Scenario A, we might have an increase in prepayments. However, we also might see prepayments level off a little bit because the people who were going to refinance that were really trying to get a better deal. They have already done so at a rate that's close to what's available in Scenario 4. It doesn't make sense for them to repeat that. Those people are already removed from our mortgage-backed security. That's called burnout, and we'll come back to that later.

In Scenario B, on the other hand, rates just increase a small amount, and there's widespread panic amongst the borrowers. We get a very high level of refinancing because people who were waiting the market out to see if interest rates would drop lower see that their opportunity is vanishing. They just run right out to refinance. Since the mortgage market is driven by supply and demand, the increasing demand might even cause a small additional increase in mortgage rates that would motivate the people that obtained mortgages at the beginning of the curve even more. Most of our prepayment models aren't this sophisticated. Most of our models reflect a higher rate of prepayments under Scenario A at 0.4 than they would under Scenario B at 0.4 because interest rates are lower under Scenario A.

Chart 3 is another example that involves burnout. This is another interest rate scenario. Let's assume that interest rates go down at first in Periods 1, 2 and 3. They increase a little bit in Period 4. We know from our last example we will get a rush of prepayments at Period 4.

They're at the same level at Period 5 and at Period 3. Then they decreased slightly in Periods 6 and 7. At Period 5, we'll have lower prepayments than Period 3 because the people have already had an opportunity to refinance at those rates that were available at Period 3. They will have already had an opportunity to refinance and chosen not to do so because they don't have the ability or initiative. If they're still in our security, if they're still in our mortgage pool, they chose not to do so, so they're less likely to prepay at Period 5.

That's called burnout. In most pools, at T equals 7, we'll have lower prepayments than at T equals 6 because the rates are level, and people will have had an opportunity to refinance at Period 6, and they will have chosen not to do so. Period 7 will have a lower level of prepayments than Period 6. They will have already had the ability, and those with the ability and the incentive to refinance will already have done so. Many of our prepayment models will have equal prepayments being equal at Period 3 and at Period 5 and also at Period 6 and at Period 7. Our models might not always move us in the right direction. There are a lot of prepayment models out there. There are commercial mortgages that try to incorporate all the factors like burnout and borrower behavior.

All the prepayment models out there do have one thing in common. Regardless of how sophisticated they are, they're all wrong. None of them will correctly predict the actual level of prepayments in the future. So we've got an important assumption in our models that we know is wrong. We could try to move to a more complex model. Many of us have probably been approached or offered fairly complex and sophisticated prepayment models that will fit right in seamlessly with our cash-flow testing software. The question is whether the prepayment, that model that we have now, is oversimplified or would it be adding too much complexity to move to other models? What's the right way to go? There is one way we can strike a balance. To answer the question as to how important the prepayment assumption is in our models in the first place is to just do some sensitivity testing with prepayments.

We can use sensitivity testing to try to strike a balance and see if we need or should even consider a more complex prepayment model. One of the obvious things that prepayments drive in our models is reinvestment. When we get prepayments we need to reinvest the money in something else under the current interest rate scenario, which in our modeling, might be different than the interest rate scenario we started off with. As I pointed out earlier in our discussion about validation, another consideration with prepayments is that the yield and the realized product spread on assets with embedded prepayment options can depend on the prepayment assumption that's used.

An example of a fairly simple prepayment sensitivity test is shown in Chart 4. The base prepayment assumption I'm using for the model is shown at the center line on the graph. The horizontal or the x-axis represents a shift in interest rates, in other words, the difference between the rate available on a new mortgage and the rate currently being paid by the borrower. A shift to zero would mean that somebody's interest rate would remain the same. A shift of 100 would mean an increase of l% in the mortgage yield curve or a shift in the interest rate that is available for them to use to refinance. The vertical for y-axis is the conditional prepayment rate which actuaries would just call an effective rate of prepayments as a percentage. If I read up the scale, my center curve line would suggest that my effective rate, if there were no changes in interest rates, is about 15% a year. The mortgage would have an average lifetime of between six and seven years or so.

In addition to our base curve, there are three additional prepayment assumptions up here. The bottom line is always lower than our base assumption. There's Alternative A where the upper line is always higher than the base assumption. There's the fourth Alternative C, which is higher than our base assumption when interest rates decrease and then lower than our base assumption when interest rates increase. In other words, it's just more sensitive all around or more conservative than our baseline assumption. If I run my model with a base curve and get results, and then I run it again with these three alternatives, I can see how much the results of my model change. I also get a feel for what happens when my prepayment assumptions are incorrect.

Another consideration related to embedded options is related to credit risk. Most of the time we don't consider credit risk to be related to embedded options at all, but many of the commercial mortgages we might have in our portfolios have make-whole provisions. We don't even consider that, or we may not even consider that they have embedded options. The borrower could prepay, but it's market-value adjusted to the current interest rate scenarios. We might not care about it. We might think of credit risk and defaults as being both beyond the control of us and the borrower. So there's no optionality involved, but the borrower of a mortgage effectively has a put option to sell the underlying property back to the lender for the value of the outstanding mortgage balance.

If a mortgage defaults, usually the value of the underlying piece of property has declined. There might be a vacancy that creates cash flow problems for the borrower. Perhaps they're not collecting the rent, so they can't make the mortgage payments, and that contributes to the default. Or maybe the value of the property has declined due to what has happened in the local real estate market in the area or something specific to that property. We don't get many mortgage defaults where the property can actually just be turned around and sold for more than the value of the outstanding mortgage loan. That's another way of saying the borrower has the option to sell a lender the property for the value of the outstanding mortgage debt. It's just like any other option. It just gets exercised when it's in the money.

A company that has issued corporate bonds has the right to sell the company to the bondholders for the amount of the outstanding debt. Corporate bonds usually have covenants that limit when that's done to bankruptcy cases.

Another important consideration in modeling investments with embedded options is volatility. Interest rate volatility and equity volatility affect the market value of assets with embedded options. If we do projections and we do testing and modeling with stochastic scenarios, the volatility assumptions are used to create and project the scenarios that we use. I want to make a distinction between historical and implied volatility. Implied volatility is used for pricing options. It might be different than, although it's highly correlated with, historical volatility, which represents the standard deviation of the interest rate or the equity index that we're talking about here.

We know how historical volatility affects our model. If we model with stochastic interest rate scenarios, the key assumption for projecting the scenarios is what's the standard deviation of the change in interest rates or change in equity in my future scenarios? Let's take an example of how implied volatility affects our model. If I have a callable bond, the issuer has an option to call the bond. As the owners of the bond, we have a short call option or a short position in a call option. If the volatility of interest rates increases, interest rates could go in either direction. The volatility doesn't specify which way they're going; it just specifies that they're more volatile. The option that the issuer owns gets more valuable because they're more likely to decrease by a large enough amount so that the borrower's call option is in the money. The call option is more valuable to the issuer of the bond, but we've got a short position in it. So the market value of our callable bond will decrease if interest rate volatility increases.

It's the same situation for a CMO. If the volatility of interest rates increases, the value of the option to the borrowers increases, but the value to the market value of the CMO that we hold decreases. Some models use an approximation to take the market value of an asset like this.

They use a single spread or single option-adjusted spread to calculate the market value of bonds with embedded options. This is usually a reasonable approximation as long as their portfolio doesn't have too many of these securities in it. It's equivalent to assuming that implied interest rate volatility is constant. Equity volatility is very similar. Implied equity volatility affects option price as well. Historical volatility is used as the standard deviation to project our scenarios.

There are three main ways of modeling volatility. The most common is to assume that it's constant. Let's say volatility is a real critical assumption for our model. For example, equity volatility is usually a very critical assumption for modeling equity-indexed annuities and supporting options. As such, we might want to consider doing some sensitivity testing instead of just using it as a constant. There are also stochastic volatility models out there, and there are regime-switching models that assume that the historical volatility that's used to project my scenario can change between two different underlying distributions or two or more underlying different values.

Some considerations for our scenario-based modeling of assets with embedded options are the type of equity and interest rate models we want to use. The broad types are risk-neutral and realistic models. I'll talk about the differences a little more, but, in general, risk-neutral models are used for pricing, and realistic models are used for most types of modeling. Another consideration applies to our stochastic or even our deterministic scenarios that we use in our modeling. The question is whether to normalize. There's not a right or wrong answer. It's just an important consideration. As an example, the U.S. Treasury curve is very steep right now. A couple of days ago, the three-month rate was 1.65%, and the long-term rate, which is now the 20-year rate, was 5.2%. That's a difference of about 3.5%. A more typical historical difference might be between 1.5% and 2%. If we think that steepness is temporary, should we revert to a more normal slope in the scenarios that we use for modeling? Again, there's not a right or wrong answer to that, but it can have big implications for our model. If we're using interest rate derivatives to hedge interest rate risk, normalization would have a big impact for existing assets in the portfolio and also for the prices of the derivative assets that are purchased as part of a future investment strategy. If I'm trying to hedge upward rates with an interest rate cap, a steep yield curve means that the interest rate caps are very expensive because the implied forward rates in the yield curve are increasing. In other words, the expectation that's built into a steep yield curve is that interest rates are going to increase. Interest rate caps are expensive under it. So if I rely on using interest rate caps to manage interest rate risk, and I assume the yield curve flattens out and normalizes, I'm going to be buying caps cheaper. It's a little bit of an aggressive assumption in my model.

Risk-neutral models are used mainly for pricing or determining the market value of assets. Risk-neutral scenarios or models get results consistent with the way an asset would be priced on Wall Street. The underlying assumption is that all investors are risk-neutral. Now I've just kind of violated that sacred law of elementary school that says, when you define a word, you can't use the word you're defining in the definition. What does risk-neutral mean? I can start by explaining what it doesn't mean. If an investor is not risk-neutral, they're risk-averse. A risk-averse investor has two available investments that they can choose between. They have two available investments with different levels of risk. They demand a higher return for the riskier investment. That's the way people are in the real world. They want to be compensated for taking risks.

In a risk-neutral world, if there are two investments with different expected returns, all the risk-neutral investors would just buy the one with the higher expected return. There's no regard for risk. We all know what the laws of supply and demand cause if we're just buying one asset and not another. The price or the asset with a higher return would be driven up. As a result, the expected returns on the two assets would be equal. In a risk-neutral world, all investments have the same return as a risk-free investment. In real life, investors want higher returns for taking risk. The amount of additional return that's necessary for taking risk might differ from one investor to another. We all want some additional return for taking on additional risk.

The reason we use risk-neutral models for pricing is because there's a large body of work that demonstrates that the price of assets doesn't depend on the risk tolerance of investors. So different investors have different risk tolerance, but asset prices reach an equilibrium anyhow. Being risk-neutral is an extreme risk tolerance, but the prices of assets don't depend on risk tolerance. It makes the mathematics a lot easier if we assume investors are risk-neutral. The math is much easier in a risk-neutral world. So we don't use risk-neutral models because the assumption that investors are risk-neutral is reasonable. In fact, the assumptions are unreasonable, but the reason we use them is because it makes the math easier. It gives us the same answer anyway.

Realistic models are used to determine the range of results of a portfolio of assets and liabilities using historically based parameters. They include a risk premium for riskier investments. Some realistic models allow arbitrage opportunities. They allow some opportunities for riskless profits, and sometimes we might care about that, and sometimes we might not. Say I have a generated interest rate that projects interest rates at quarterly or monthly time intervals. It might assume at a level in between the quarterly or monthly time intervals, that there's some arbitrage opportunities available by buying and selling securities that mature at different points in time within the month or within the quarter. I really wouldn't care about that in my modeling because I know that arbitrage opportunity might exist. I wouldn't exploit it in the model or within the strategies of my model.

A different example would be if I was using an interest rate generator that generated a yield curve that was very steeply inverted. Perhaps the long-term rate is much less than the short-term rate, and it was so inverted that some of the implied forward rates that were built into the yield curve are negative. That gives me an opportunity for riskless profit. If that situation existed, people could put money in a safe and make money on it. That's something I'd want to avoid because that's an arbitrage opportunity that we could exploit in our modeling, even accidentally, without realizing it.

There is another consideration for assets with embedded options. As challenging as it is to model the asset side of the balance sheet, we need to look at the liability side, too. As an example, let's consider a single premium deferred annuity with a portfolio interest rate crediting strategy. We take the portfolio rate, we subtract a spread, and we get a credited rate. The credited rate is reset annually. Let's say the company has a few interest rate derivatives. Let's say there's some interest rate caps that pay off if interest rates increase to hedge the risk of increasing interest rates. Let's say the company has a minimum guarantee credited rate on the single premium deferred annuity (SPDA).

They also buy some interest rate floors that pay off if interest rates fall. They also use some interest rate swaps to convert some of their fixed debt to floating rate debt, which effectively lowers the duration of the assets in the portfolio. There is an investment policy question that's very important to reflect in our model and to reflect in practice. I think it's more common to get it right in practice and miss something in our modeling.

Since we have portfolio interest rate crediting, we take the portfolio rate, and we subtract the spread to get the credited rate. The questions is, which, if any, of those derivatives — the caps, floors, and the swaps — get their investment income or their income included in the calculation of the portfolio rate? With the floors, it might not make too much difference. If the floors are paying off, it means that interest rates are low. The company is probably not able to earn their full spread on a product because the interest rates are so low. The rate that the policyholder gets might still be the minimum credited rate regardless of whether the income from the floors is included in the calculation of the portfolio rate or not. But, in a rising interest rate scenario, the question is, does the company just keep the investment income from the caps and the swaps to offset their risk, and to offset the fact that they might get some disintermediation? They'd have to sell some assets in a capital loss type of situation. Does the income from the caps and the swaps get included in the portfolio rate and passed onto the policyholder to give them a more competitive interest rate, given the rising interest rate scenario? It's really important to consider the intention of the company in our asset modeling.

For stochastic modeling of equity-based products, our main concern is the two types of volatility we talked about previously and the expected return. One of the questions for expected return is what level to set it at. The other question is whether to incorporate some type of mean reversion in the returns that are created. Colin Devine spoke about this at the General Session in a GAAP context. He said a lot of companies are up against a collar as far as their DAC unlocking. Since the market has declined in the past few years, mean reversion or the returns suggest that the expected return for the equity markets going forward is about -14%.

Although there are good arguments on both sides, I personally don't like using mean reversion for equity returns for a couple of reasons. There have been extended bull and bear markets in history, and these extended bull and bear markets are more likely to occur and actually happen in some of our stochastic scenarios that we model with if we don't put constraints like that on our equity projections. There's also the efficient market hypothesis, which does have some empirical studies to back it up. It says that all information that's incorporated in historical equity prices is fully incorporated in current equity prices. Looking at the past has no predictive value for us for the future, but mean reversion would suggest the opposite.

I'm going to begin the last section of the talk, which is about using models to assist us with portfolio management. One of the main tools available to us are stochastic asset and liability models, and I've talked about those a little bit already. The main strength of stochastic models is the ability to demonstrate the interaction between assets and liabilities. By contrast, if we're just looking at the asset side of the balance sheet, there are a lot of linear programming-based asset optimization models out there that we could use to obtain an optimal asset portfolio that satisfies a given set of constraints. Stochastic modeling adds value to that because it can incorporate both sides of the balance sheet, and it incorporates the interaction between assets and liabilities.

If we want to improve spreads from our portfolio, usually we need to take on some added risk to do that. We can use models to decide how much risk we're willing to take on. One of the opportunities for an insurance company to increase asset yield and, therefore, increase spread is credit risk. One of the key modeling issues with regards to credit risk is that there are two types of credit risk. There's a risk that an asset, a bond, or a mortgage can default, and there's risk that it could be downgraded. Downgrades might not cost us anything because if we hold the bond to maturity, and the issuer makes all the payments, a downgrade would turn out to be irrelevant. There are three reasons why downgrades can matter. First would be a case where a bond's downgraded by the NAIC Securities Valuation Office to a higher numbered NAIC class. Then we need to hold more risk-based capital. We need to hold more capital in our company to support it.

The second case would be where a company reevaluates whether to hold the bond or not. They might sell some of the bonds that get downgraded at a higher discount spread, which does cost us money. A third reason is if a bond gets downgraded, the probability of an actual default increases exponentially if it gets downgraded from one class to another. If we're going to use our models to consider changing our company's exposure to credit risk, the two key types of credit risk that our model should reflect are defaults and downgrades. Another key consideration for this type of model is that to really get a feel for the distribution of possible outcomes, we would want to model credit risk stochastically instead of using a simplified assumption that a fixed amount of each asset defaults. That is something we might use for a cash-flow testing model.

Another key consideration is that credit spreads and bonds are tied to expected rates of default. Credit spreads on bonds are fairly high right now relative to their historical levels but expected defaults are higher than their historical levels also. If we do think that the credit spreads on bonds reflect things that have happened in the past like Enron, Global Crossing, World Com, and their compensation for those events, then we might think that that's an opportunity for us. A final key consideration to modeling of credit risk is that defaults are not independent. They're correlated to the economy, and there's also correlation between bonds within the various sectors of the bond market, and there's also correlation with real estate and mortgage holdings that are in the same geographical area.

Another opportunity for companies to increase their spread is by taking duration risk. A good number of insurance companies already do that. They have an asset portfolio with a duration that's higher than their liability portfolio. If I look at the last 50 years of history, the average one-year Treasury rate is 5.94%. The average three-year Treasury rate is 6.34%. It's 40 basis points higher. If I had a yield curve that looked like this, and I had a one-year liability, I could invest the money in a one-year bond. As a different alternative, I could also invest the money in a three-year bond and sell the three-year bond after one year to retire my one-year liability.

Let's say I bought the three-year bond, and interest rates didn't move at all. If the yield curve looked like this historical average, I'd get an additional 40 basis points, but I'd also get a capital gain because when I sold the bond, after one year, the three-year bond would now be a two-year bond. Since my yield curve is upward sloping, the two-year rate is less than the three-year rate, so I would also get a capital gain when I sold the bond, and that would amount to another 50 basis points or so with this historical average yield curve. So I can pick up 90 basis points or almost a full percentage point by taking on a very small amount of duration risk.

There are two theories about the shape of the yield curve. There's the expectations hypothesis that says an upward sloping yield curve means that interest rates are going to increase, and the amount of the increase is built into the implied forward rates of the yield curve. The hypothesis would say that buying the one-year bond and the three-year bond would be equivalent on average. In the case where I bought the three-year bond, interest rates would increase so that I'd get a capital loss that would offset the additional 40 basis points that I get in yield. The second theory about the shape of the yield curve is the liquidity preference theory that says most of the time the yield curve slopes upward in order to compensate investors who take on the added risk of investing further out on the curve.

Empirically, based on past data, the liquidity preference theory, turns out to be correct more often. About 60% of the time I'm better off buying this three-year bond than I am buying a one-year bond. The average amount that I pick up is 92 basis points or, again, almost a full percentage point. There is risk in this approach. The standard deviation of my outcomes buying a three-year bond is about 3.5% of my principal. The worst case, historically, would be if you did this in July of 1980. Between July of 1980 and July of 1981, interest rates increased 6% or so on the short end of the curve. That would give me a loss of about 12% or 1,200 basis points. This is a pretty simple example, and something that's rated PG-13 compared to most of our asset and liability models. To really determine what level of duration risk is appropriate, we need to consider using a tool like a stochastic asset/liability model.

I wanted to wrap up by discussing a couple of statistical tools that we can use with our models and also talk briefly about benchmarking, which is not a statistical tool. It's a different technique or different approach all together for improving our yields. The first statistical measure I want to talk about is value-at-risk. What the statistical measure value-at-risk means is that it's a maximum loss at a given confidence level over a defined period. To put it another way, it's a percentile of the distribution of the bad incomes or negative results. There can be a great deal of work involved with analysis of this type, and a full discussion is beyond what I wanted to talk about here. What I wanted to point out was that if we get results from our stochastic analysis, we can use value-at-risk as a statistical measure of the risk.

We typically look at standard deviation of the results as a risk measure. If I had a two strategies, for example, and one strategy just increased the positive results and left the negative ones largely unchanged, it's going to have higher standard deviations. We might not care if it just improves some results and leaves the other ones the same. We might not care. Two statistical measures, the value-at-risk and the contingent tail expectation (CTE) would pick that up. We'd see different values for the value-at-risk and a CTE. We'd see a higher expected return and the same level of risk so we'd know which strategy is superior.

Contingent tail expectation just means the average result for a selected group of worst performing scenarios. The simplest case would be CTE zero. That would mean it's the average result for the 100% worst scenarios or, in other words, just the average result for all the scenarios or my mean result or, in other words, the mean. CTE zero is the same as the mean. I have to subtract the 75 from a hundred. CTE 75 means the average result of the worst 25% of the scenarios, and CTE 90 means the average result of the worst 10% of the scenarios. CTE is used in Canadian valuation requirements for annuities with guaranteed living benefits. It's starting to catch on in the U.S. as just a general risk statistical measure that can be used.

Benchmarking is a totally different approach to portfolio management all together that really just involves passing the buck. The responsibility for risk is delegated to the portfolio manager. A hypothetical benchmark portfolio is set up. The portfolio manager can buy that benchmark portfolio, but typically they're going to buy a different portfolio because he or she is evaluated based on the performance of the actual portfolio they buy to the benchmark portfolio. The technique is used very commonly in banking and pension investment, and there's starting to be some interest for people using it for life insurance portfolio management. A typical benchmark portfolio for insurance would be a liability duration-matched portfolio of investment grade corporate bonds.

Another special case would be a portfolio of bonds, calls, and put options that exactly replicates the cash flow from the liability portfolio. That's a special case of benchmarking that some of us call transfer pricing. The manager has the authority to deviate from the benchmark portfolio within certain risk limits and within certain regulatory and rating agency considerations. Of course, the manager is evaluated on the performance of the portfolio relative to the benchmark portfolio.

FROM THE FLOOR: I have a question on floating rate assets.

MR. HOUGHTON: Yes.

FROM THE FLOOR: For those assets, the market value is a function of spreads as they move, and I haven't seen any generators yet which give you stochastic spreads. For most of our models, the Treasury rate, rather than the spreads, moves on a stochastic basis. What do you do there?

MR. HOUGHTON: Yes, these are simplified approaches. Something to consider is that credit spreads are fairly high relative to the history, and what's tied to that is expected default rates are higher, too. So we can use default rates that are higher than what we've seen in history using a Moody's study of the past 30 years. Pulling our default rates out of that might not be appropriate because they're tied together. The spreads are high, but the expected defaults are high. What we can do is use current spreads that are high but also increase our expected level of defaults. We'd have at least some consistency that way. Another approach would be to grade down the spreads to what we think is an appropriate long-term level. Also, start with higher expected default rates and grade those down to what we think is an appropriate long-term level. That's another simplified approach. As you mentioned, a third approach could be creating a spread stochastically, although that does add a lot of complexity. Depending on how many of those assets you're modeling, we'd have to really consider whether adding the complexity is worth it or not.

The question from the floor is when would I use CTE and when would I use value at risk? What are the advantages of each? They are just statistical measures. I would almost always use both of them. The amount of work required to calculate those is very small considering how much work I've done to create a stochastic asset/liability model and run a stochastic asset/liability analysis. I guess my answer to that is I would always use both. Typically, we'd see value-at-risk at the 90th and 95th percentiles. In my analysis, I usually create a table that has CTE 90 through CTE 50. It just gives me a quick way to compare one strategy to another. It's something you can calculate in a couple of minutes in a spreadsheet. They're both statistical measures. They're easy to do. My answer would be I always use both.

FROM THE FLOOR: What do you use for developing assumptions? When do you use spread-based or more sophisticated models?

MR. HOUGHTON: I use both depending on the client and what type of securities it has and how much exposure they have to those securities. There is, for government mortgage-backed securities, various historical prepayment information available from the government issues. Ginny Mae, Fannie Mae, and Freddie Mac all have a lot of statistical data on past prepayment rates available on their websites. That's certainly something I would use to develop an assumption for a certain client or for a certain company if they had those securities in their portfolio. Typically, if an insurer's more risk averse and just has a small amount of securities in its portfolio, I would use one of the simpler models that I talk about that would have a spread-based prepayment model that has a change in interest rates and then a specific prepayment rate that would occur under that change in interest rates for a small amount. If it was more sophisticated, there are commercial models out there. I don't want to endorse any of them, but Andrew Davidson and Bloomberg and some of the investment banks have commercial prepayment models. I occasionally use them, but the one thing I don't like about them is that they're sometimes proprietary. They're a black box. So it's difficult to see what assumptions are being used other than maybe a broad general description from the manufacturers. What I do is try to validate it and run a mortgage-backed security through a lot of different interest rate scenarios and see what prepayment level comes out. If I run it through some parallel shifts, I can run a prepayment model through an immediate decrease of 2%. Then, interest rates stay level after that. I can reverse engineer what the model is coming out with by just running it at different levels of fixed prepayment rates until I get the same pattern of prepayments. I could see that under that scenario I'd get a 30% CPR, and that seems reasonable. I use both commercial models and more simplified interest rate models.

To answer the question, I have not seen any models with stochastic default rates. I've seen defaults modeled stochastically using a fixed rate of incidence.

Another question from the floor is whether there is any software that models collateralized bond obligations (CBOs). CBOs are grouped with collateralized loan obligations. I'm not aware of software that models those well. I've seen companies with mainly the individual issues, so it does seem like an open market.

CHART 1
Price Curve Example

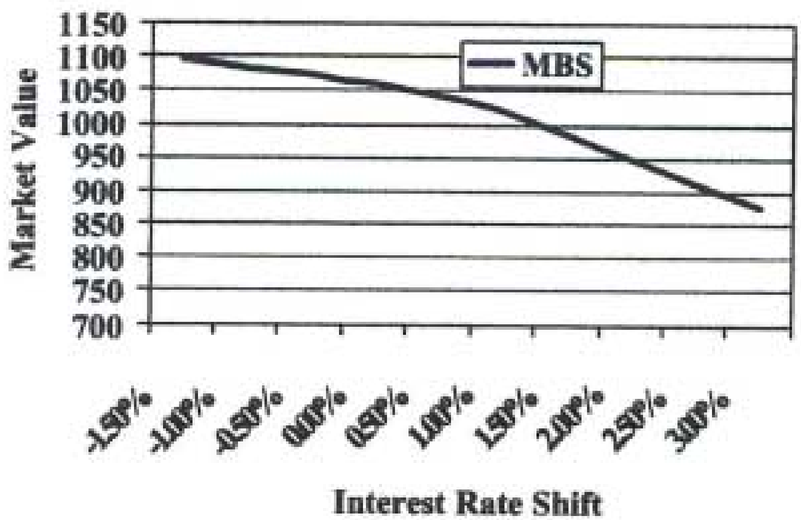

CHART 2
Investor Behavior in Prepayments

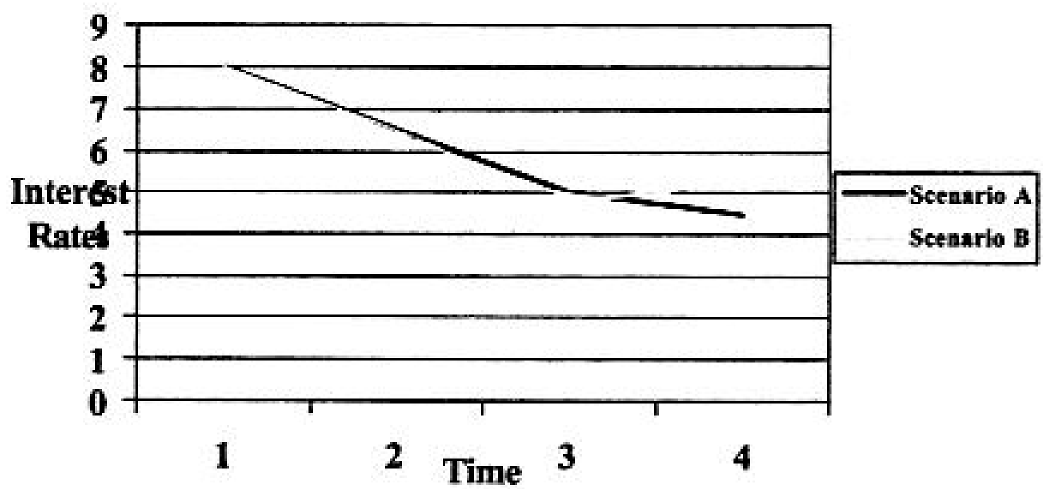

CHART 3
Burnout in Prepayments

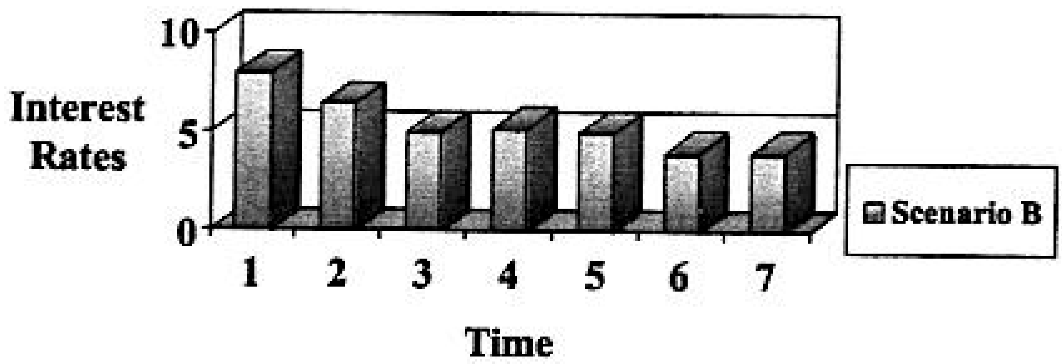

CHART 4
Prepayment Curves