1996 VALUATION ACTUARY
SYMPOSIUM PROCEEDINGS

SESSION 27

Asset Modeling I

John E. Heinmiller, Moderator

Peter Fitton

Joseph M. Rafson

ASSET MODELING I

MR. JOHN E. HEINMILLER: The material to be covered in both Asset Modeling sessions will be as follows: I will begin by talking about the purposes of building an asset model and the resource requirements for building a model. Joe Rafson will then talk about model validation and model granularity and understanding spreads. Peter Fitton will finish this session by talking about model interest rates.

In the second session, Asset Modeling II, Peter Fitton will cover modeling securities with options and derivatives. I will cover the topics of corporate government debt, mortgage-backed securities and collateralized mortgage obligations (CMOs). Joe Rafson will finish up talking about real estate, equities and investment strategies.

First, I want to introduce our speakers. Peter Fitton has been employed with SS&C Technologies for three years and is a subject matter expert in advanced interest rate modeling and fixed-income securities. Peter holds a master's degree in mathematics from the University of Waterloo.

I am the director of development of analytics products in the Chalke Division of SS&C Technologies. I have been with Chalke for 12 years.

Joe Rafson is a Fellow of the Society of Actuaries and a consultant in the Chicago office of KPMG. He specializes in asset/liability management issues. Joe Rafson performs cash-flow testing for several companies each year. Each year he reviews about 15 annual memoranda. Joe is a member of the Finance Practice Advisory Committee and is a frequent speaker at Society meetings.

We are going to begin our discussion of asset modeling by exploring the various purposes for creating an asset model. It is critical to identify, at the beginning of the model construction phase, the purpose of building an asset model and the criteria that can be used to measure the quality of the model. It is common to find model builders lured into the details of particular assets in the portfolio, and forget what the motivations are for creating the model.

There is some protection from falling into this common trap by initially constructing a crude model that captures common attributes of securities within the portfolio, and then using what time remains to iteratively refine the model's details.

The most common purposes for building an asset model include the following: valuation, risk management, performance measurement, rating agency surveys, spread management, and cash-flow modeling. Most actuarial model building has centered around product pricing, so it is common to find actuaries trying to fine-tune their asset model by comparing their model prices to prices provided by the firm's portfolio manager. Except for riskless bonds, trying to match prices with a model can be frustrating because of lack of information on how securities are priced.

A purpose for asset model building that has increased in importance in the last few years has been the rating agency surveys. A.M. Best and Standard & Poor's are both taking a closer look at portfolio volatility. The reporting requirements for A.M. Best are more severe, requiring option-adjusted analysis on nine interest rate shocks.

Another purpose for building an asset model is performance measurement. Performance measurement can mean different things to different people. What I mean by performance measurement is really a measurement of the portfolio manager. It comes as a great surprise to many actuaries that a significant portion of compensation for the firm's portfolio manager is calculated based on some assessment of the performance of the assets under management. So it is critical to review how this performance is measured and to ask if the benchmarks established for the asset manager are appropriate. These performance benchmarks and the constraints placed on the portfolio managers are closely watched by Wall Street, and have significant impact on the creation of new security types.

A good study of the impact of performance measurements on the securities industry is a study of the emergence of the collateralized mortgage obligation. I recall the first company that called our firm inquiring whether we could model a CMO. This company had recently purchased these new securities, but could not begin to describe the cash-flow characteristics.

Now I'll turn the discussion over to Joe Rafson who will cover model valuation and spreads.

MR. JOSEPH M. RAFSON: I'd like to say a few general words about what I do. Since I review so many actuarial memoranda as part of the statutory audit process, I perhaps have a different view of things than most people. Many of my clients have excellent quality memoranda, but I am frequently also given the task of reviewing some of the more problematic ones. So my view of the world is a little bit tainted by what you're not going to see at the Valuation Actuary Symposium; that is, some of the industry leaders and some of the best practices.

Model validation can be one of the most time consuming tasks in cash-flow testing. The choice of items to validate can be somewhat problematic. The more items you want to validate, the harder your validation becomes. The purposes of your model become very important. How accurate do you need to be? Are you doing a basic sensitivity test, or are you doing a forecast upon which management compensation may be tied? It may be very important to your choice or your threshold for validation. What is close enough? Normally, you'll have a baseline for your liability or asset side which ties out pretty much to the dollar. For instance, if you're inputting the book value of assets, you should be tying out to the book value of assets. Normally, on the liability side, you're starting with either statutory reserves or account value. Those should tie out pretty directly.

The next corresponding measure becomes a validation item. On the liability side, statutory reserves become a validating item if you've started with account value, similar to account value if you've started with statutory reserves. Face value might be especially important for term insurance. Annualized premium in force might be especially important for accident and health (A&H) coverages. For asset modeling, the market value of assets may be your validation item if you start with book value as your baseline. I'm assuming the market value of assets is being produced by the software you're using. So what could make your market value be off are problems with the calculations, problems with your data, problems with your assumptions (particularly your spread assumptions), and problems with your interest rate inputs. So if any of those items are off, you're not going to validate to the market value of your assets. If the market value of your assets is wrong, and you have to sell assets, your cash flows are off. Then it is more than an indicator value. Asset validation is very useful in checking some of the other inputs into your test. If you can break down your validation by asset class, you can narrow down any problems in performing a validation. Let's discuss some of the experience items. Now, we're moving to the income statement rather than balance sheet items to validate.

On the liability side, how good a predictor is your model of lapses, mortality, expenses, and premium income? On the asset side, investment income and yields will be the items to validate. It's a red flag, and a big red flag, if you're not agreeing to aggregate portfolio yields or investment income. We read memoranda where more income is produced by the assets in a model than are produced by the assets of the company, and that's a problem. Yields can be fairly easy to validate since most of the available software has the yield right there at the bottom of the page. There are always reasons not to tie out exactly. You're not going to be modeling all of your assets in a cash-flow test. If you're excluding your real estate, your common stock, or portions of a commercial mortgage or bond portfolio, it's a fairly simple task to back out the yields on those assets to get an overall feel for the yield you should be getting on what is modeled on both a book and a market-value basis. If you're not tying out well, that could be indicative of a problem. If you're 50 basis points off of a reasonable estimate, it can have a dramatic effect on your projected earnings, particularly since many models are built with portfolio crediting strategies. So any difference of what you're earning on your assets goes straight to the bottom line.

There are other ways to validate what I call a year-to-year validation. I ask the question, how good was last year's model in predicting what happened during the year? The specific answers to this question provide a real opportunity for model refinement. The better the feedback loops you can put into the process, the better your entire model will be, and the more useful it will be. If you're using the same model or a similar model with known differences for different purposes, use the same validation steps for each model. It's simply not worth going through all this effort for a model that's not a good predictor, unless you view the tests only as a way to get the regulators off your back. Then, perhaps, unfortunately you may have a lower threshold for validation.

In terms of granularity, we're talking about how much we can take our individual assets and simplify their modeling. I'm told there was an asset/liability management seminar in Reno the other day. The panel felt it was okay to model assets into similar groupings, for whatever reason. Yet, for most bond and commercial mortgage portfolios, seriatim modeling is, by far, the most common in the industry. The model builder programs on the liability side are better than on the asset side, so it becomes a more difficult task on the asset side to get a good fit. In addition, the uses of the model will affect the degree of granularity you need. Some of the reasons for granularity or aggregation of assets could be time considerations or data considerations. I think those reasons are becoming less and less a factor as the software gets faster and the computers are much better. I'm sure we all remember the days when the cash-flow testing models used to have to run over the weekends. Now it's much more common to have a model running over lunch time. I think it's a much more acceptable process to have a seriatim listing; it's easier and it's more precise.

One of the things I use as a red flag is when people start aggregating assets that aren't normally aggregating. In one case, seeing the phrase *aggregate asset* in a cash-flow memorandum was the red flag that eventually led us down a path to not issuing an unqualified opinion on the entire company. When equity is modeled it's almost always boiled down to one asset. When real estate is modeled, it's frequently properties that are combined or cash flows that are combined into a one-line input. But for the bulk of the assets in the industry, like bonds, you typically see a seriatim listing. It's not that aggregating assets is a problem inherently; it's just one step that can cause errors.

When you're talking about liability granularity, the software is much better. Typically, a company will model a large portion of their liabilities and somehow map the nonmodeled portions into those modeled liabilities. We like to see the mapping of nonmodeled plans into similarly modeled plans, rather than grossing up an entire liability model. We feel that's more accurate.

Sometimes companies have a 1,000-cell model, as if the model will validate well just because they have broken it into so many pieces. My only comment there is, obviously, the model is only as good as the individual cells, and a 1,000-cell model has the opportunity to be more accurate, but also has more opportunities for error in it as well.

Understanding spreads. Spreads are one of the most important assumptions that will go into your model. You typically have projection periods of 20 or 30 years. A typical portfolio is going to begin rolling over a large portion of its assets immediately, and by five to seven years, you'll probably have the majority of assets in reinvested assets. The spreads that those assets are purchased at will be extremely important to your model results. We know that spreads are set by a complex marketplace of issuers, who are borrowing money, and the purchasers of assets. It's a supply-and-demand situation. Whatever you think a seven-year A-rated noncallable bond is going to yield is sort of irrelevant to what you can actually find in the marketplace. Obviously, you still have to set assumptions.

The assumptions on asset returns are typically of two types. One is treasury yields plus a spread, and another is a multiple of a treasury yield plus a spread. If it's just treasuries plus a spread, you're using some kind of an average spread, typically on an ultimate basis. If you use a multiple of the treasury plus a spread, you're recognizing the fact that spreads tend to widen when interest rates are higher and tend to narrow when interest rates are lower. The assumptions, the multiples and spreads, are typically set using historical input into a regression analysis. Obviously, regression is not a guarantee of what is going to happen. That sort of relationship may not hold in the future, and one of the sensitivity tests that we very rarely see is the test of what happens if desired spreads just aren't available. If we can't achieve the asset returns we think we're going to achieve, does this cause us a severe problem? What are we going to do in that case?

Spreads are also one of the areas where we see actuaries manipulate their models by adjusting spreads. If I can just milk an extra ten basis points out of a model by assuming an extra ten-basis-point spread on an average asset, I'm going to be taking home ten basis points more of income, especially if I'm tying my competitor rate to the treasury rate underlying those spreads.

What's conservative in spreads is not clear. You get a different answer when you're selling assets than when you're buying assets. When you're selling assets a wide spread is conservative and when you're buying assets a narrow spread is conservative. And at some companies, we actually see separately modeled investment classes for existing assets and purchased assets to allow for conservatism when buying or selling assets.

Market value of assets, as of the beginning of the test, is a source of validation of your initial spreads. Most of the commercial software systems allow you to vary spreads between initial spreads and ultimate spreads. If you are validating the model, and spreads are narrower than historical averages, you can grade to ultimate spreads with most of the software.

My last comment on spreads centers on the default spread assumptions in both the TAS and PTS systems. These defaults are not at all assumptions. But what's a little bit peculiar, though, is how commonly those defaults are used without changes. Perhaps it's an indicator that not all actuaries pay enough attention to spread assumptions. Sometimes spreads can tell us things. I have one client who had an abnormally high spread between yield on public and private placement bonds of a similar quality. He was questioned about whether this was accurate. The answer came back that the underwriting of the private placements of that company was somewhat lax. So what was labeled as an A-rated private placement maybe was not quite as high a quality as an A-rated public bond; in fact, it was an accurate set of assumptions.

MR. PETER FITTON: My topic area is going to be modeling of interest rates. My intention is to introduce the concept of interest rate modeling. More specifically, I want to explain the important concept of including spreads for risk in an interest rate model.

We'll start by talking about a spread for interest rate risk, which is the most important and probably the least understood. The spread for interest rate risk leads to two different model forms: first, there is what is called the risk-neutral model, which is used for discounting uncertain cash flows to find a present value; second, there is the realistic model, which is used to generate scenarios for stochastic cash-flow testing. We use a realistic model when we want to randomly generate a scenario projecting some reasonable future behavior from a model, in order to see what happens. In other words, we want the realistic model to determine a probability of passing or failing a solvency test.

Let's start by thinking about an initial yield curve, as observed in the market, corresponding to the date at which you want to value some uncertain future cash flows. From that yield curve we can derive a set of forward rates. This forward rate curve is, as many of us know, not a particularly good predictor of future interest rates. However, this forward rate curve has been drawn from a set of spot rates, which we know are the correct rates for discounting future cash flows.

The spot rate curve is a set of rates that we can derive from the initial yield curve that we observe in the market, and that we can use for discounting future cash flows when those cash flows are *certain* -- we know that they're going to occur. The thing to observe about these spot rates is that they embed a premium over predictions of future interest rates. Nevertheless, we don't have to know the market's prediction of future interest rates in order to discount. All we have to know are the spot rates that we derive from an initial set of prices.

In a similar way, we can get a forward rate curve from that spot rate curve. Even though the forward rate curve is not a good predictor of future interest rates, and even though there's a spread (which we call a term premium) between it and expected rates, we can still discount future cash flows, without knowing what the term premium is. We do so by generating a set of scenarios around that forward rate curve, developing cash flows along those scenarios, and discounting along the scenarios at the short-term interest rate to the present time period.

A model that produces scenarios that are centered around a forward rate curve (derived from the initial market data) is called a risk-neutral interest-rate model. That risk-neutral model that we can use for discounting uncertain cash flows is, of course, not realistic. It's not realistic in the sense that rates coming out of the model are too high in comparison to history and in comparison to market expectations. Rates are too high by a spread equal to the term premium.

The other thing that we notice about the scenarios is that yield curve inversions happen too frequently in comparison to history. That results from the high rates, because our model structure (if it's any good) will reflect that yield curve inversions usually happen when rates are high.

So, that's why, when modeling interest rates, we must have two forms. We need a risk neutral model to find the present value of surplus; this model is unrealistic because it embeds the initial term premium into the expectations of interest rate paths. If we want to produce scenarios for testing realistic potential futures, we need a realistic model that separately identifies the term premiums. That way, the scenarios will be centered around something close to the market expectation of future rates, rather than the forward rate curve that systematically overstates market expectations of future rates.

When constructing a model of interest rates, whether it's a risk-neutral model discounting future cash flows or whether it's a realistic model for stress-testing purposes, two basic approaches have been taken historically. One is an ad hoc approach, and the other is an econometric approach. The ad hoc approach started on Wall Street, where simple risk-neutral models were used for derivatives pricing. These models were generally simple, because only the expectations of these models really mattered very much to derivative pricing. It didn't much matter whether the model had realistic behavior to price bond options or swaps or whatever needed pricing. They were typically one-factor models; that is, models with one source of random variation underlying the yield curve process. Furthermore, users of these models typically haven't worried about things like whether negative interest rates occur under the model, or whether the behavior of the model was going to be brittle because it was too sensitive to initial market prices.

However, when we produce scenarios for a realistic testing purpose, we need to be concerned about whether the model behaviors are anything like what occurred historically or what we believe might happen in the future. Thus, the approach more appropriate to the actuarial setting is an econometric approach, which involves actually studying interest rate history in order to construct a model that behaves very much like it. Of course that's much harder to do, which is why the simple models proliferated on Wall Street. Most Wall Street investment banks don't have to econometrically model interest rates for risk analysis because the banks have such short-term time horizons. Insurance companies, our employers or the companies for which we consult, have a much longer time horizon over which to be concerned about interest rate risk; the necessity of an interest rate model that produces realistic behaviors is much greater. This means that simply taking a model that was used for derivatives pricing on Wall Street a few years ago, and directly applying it to the insurance company setting, may not be the best approach.

When I mention an econometric approach, I mean an approach that looks at global behaviors of the interest rate process. One noticeable characteristic is that nominal interest rates have never been negative; in fact, they've never been extremely low (below 1%). Interest rates have never been extremely high, either; that is, much over 18%. Also, the volatility of long-term rates has been less than that of short-term rates for the historical period for which we have long-term rate data.

There is one special thing to note about one-factor models. These models have one source of random variation underlying the yield curve process, meaning that all yields (or all spot rates, or all forward rates) are perfectly correlated. Yields of every maturity always move in the same direction at the same time. This characteristic is strongly at odds with history. Historical yields have been highly positively correlated, but not perfectly so. In fact, over the last couple of years we've been in a situation where long rates moved down at the same time that short rates moved up.

There are some special considerations when using behavioral models in combination with interest rate models. On the asset side, when we talk about behavioral models, we're referring to models of default and prepayment. On the liability side, we're talking about models of lapsation and mortality. There may be some subtle dependence of mortality or morbidity on interest rates, but they are generally not modeled that way.

For a better example, let's think about a model of default. Suppose there is a model of corporate bond default built into our asset model. We may look at history and build a statistical model, perhaps a regression of default rates and recovery percentages for corporate bonds as related to historical series of interest rates. If we construct a regression equation, and thus a realistic, though simple, econometric model of the relationships between interest rates and corporate bond defaults, and then use that in calculating the present value of surplus, we will have made an error. This is a very common, frequently made error, even among many sophisticated finance people. The problem is this: just as in the Treasury yield curve, there is an embedded premium for interest rate risk, so for corporate bond yields within a particular rating class (for example, the A-rated industrials) there is a certain premium for default risk. The cash flows produced by our basic interest rate and cash-flow model will be reduced in frequency and magnitude by our default model, which says that some of the interest and principal payments will not occur. But this is not enough. If we then use the Treasury spot interest rates to discount those reduced cash flows, we will not have taken into account the extra spread demanded by the market for default risk. And so, while the corporate bond yields generated by our model will be at a certain positive spread to Treasuries, the spread is unlikely to be enough to match market yields if the default rates come from an historical regression.

The same thing is true of prepayment models, and this is where the mistake is most often made. If we are using a historically estimated prepayment model to describe how prepayment rates are related to interest rates, we need to include in our discounting procedure not only the Treasury rates, but also an extra premium for the market price of prepayment risk. This additional spread can be derived from the prices of liquid mortgage-backed securities.

On the liability side, since there is no liquid secondary market for trading insurance liabilities, and there is no market for trading mortality, it is unlikely that anyone will be able to derive a reliable estimate of the market risk premium for lapse risk or mortality risk. That means that when using an interest rate model to discount liability cash flows and obtain a "market value of liabilities," the answer is very approximate. Not only is the price of liabilities approximate, but liability duration, which depends on the difference between two prices, is also approximate (at best). If you're in a situation of calculating liability duration and comparing it to asset duration, both of which have been calculated using an interest rate model, you can put much more confidence in the asset duration than in the liability duration, assuming that you have devoted the same rigor to each. An attempt to match the two durations is unlikely to provide any more than very approximate immunization.

In summary, the discount rate for discounting an uncertain future cash flow is composed of an expected future Treasury spot interest rate plus a premium for interest rate risk (together these form the forward rate that one can derive from an observed Treasury yield curve) plus additional risk premiums for other risks such as default risk, prepayment risk, lapse risk, mortality risk, etc. These additional risk premiums can be derived from market prices of traded assets, but are much more difficult to estimate for insurance liabilities.

MR. MICHAEL P. HEALY: What about the PTS/TAS default model?

MR. RAFSON: The TAS and PTS defaults are not always appropriate for all companies, and yet they are a very common assumption used in cash-flow testing. My only comment there is that a valuation actuary should be comfortable with those defaults, and I think sometimes the steps to get comfortable with those defaults or to set other assumptions are sometimes not performed. Although I generally find the defaults to be quite good historically. The question is, is this appropriate for your investment department, and the way your portfolios are managed?