RECORD, Volume 31, Number 1*

Life Spring Meeting
New Orleans, LA
May 23-24, 2005

Session 68 SEM
Credit Risk Models and Mathematics: Part 1

Track: Pensions

Moderator: FRANCIS P. SABATINI

Panelists: GEORGE A. HOLT
ADAM C. GIRLING†

Summary: Credit risk is a significant part of an insurance company's risk exposure and has not been an area of focus for the actuarial profession. This seminar seeks to provide an overview of the approaches currently used to measure credit risk exposure within the banking and insurance industries and contrast those approaches with traditional actuarial projection techniques.

MR. FRANCIS P. SABATINI: Credit risk models and mathematics is a great topic from my perspective. I think it's a subject that actuaries need to learn more about. We have two great speakers who have a tremendous amount of expertise and experience in this subject. George Holt is managing director of consulting with Kamakura Corp. Kamakura is an interesting company. They're a risk solutions business providing software information services and consulting, and one of the things they do is credit risk modeling, among other things. George is managing director and leads their quantitative finance consulting practice. Adam Girling is a manager with Ernst & Young in the financial services advisory practice. He spends a lot of his time working on the design and implementation of credit risk management, modeling and economic capital systems. He does a lot of work with both banks and insurance companies and he's very actively involved in some major efforts in terms of Boswell two implementation with a couple of major banks.

Presentation materials can be downloaded at: http://handouts.soa.org/conted/cearchive/NewOrleans-May05/068_bk.pdf

MR. ADAM GIRLING: I'll be focusing primarily on just one aspect, but a fairly major one in the credit risk modeling world, and that is the probability of default estimate. It's something we all use in various applications. One of the things I'm going to walk through is some of the different approaches that are out there to look at default. I'll include some of the models we use, both by the vendors and some of their service offerings, and also by some institutions within their internal approaches to probability of default. The second thing I'm going to touch on briefly is the introduction to some of the different portfolio modeling considerations you have when we start translating that probability of default measure into a credit risk distribution.

To put in a little context the credit risk-modeling universe, the language around credit risk has somewhat standardized over the last decade. There are a number of fundamental building blocks within the credit risk-measurement universe. You can divide these into the input perimeters, which are somewhat counter-party specific, and then other input perimeters such as loss default estimate and exposure estimate, which are more specific to individual facilities or instruments. The probability of default that could look at the underlying components in credit risk-management, which are influencing that, probability of default largely dependent on counter-parties' financial strength and its ability to meet obligations usually independent of those obligations. We look at probability of default as it relates to an obligor and its ability to meet all its obligations whether they are bonds, repayments on a loan or as a counter-party to an insurance transaction. One of the key issues in credit risk-modeling is how we then look at correlation between probability of default estimates and the different obligors.

I'll briefly mention probability of default models. They are relatively less sophisticated. There's less data around with which to perimeter loss probability of default models, but they are more focused on some of the structural characteristics of the exposure. Many of you would have seen the rating agency statistics that demonstrate that. A senior secured bond has a higher recovery rate and a lower probability of default than, for example, a subordinated bond. The exposure default models are very simple in the case of something like a regular corporate bond. It gets far more complex when we have to look at situations where the cash flow profiles of a particular instrument are related to similar underlying risk factors that might also influence the probability of default or loss of default or other credit risk perimeters.

One of the reasons why probability of default is the focus is Boswell two, in particular. It allows internal measures or probability of default to drive regulatory capital and economic capital estimates to the banks. That brought a lot of focus on the validity of the probability of default measure. Consider its performance and stability over time. How reflective is it of actual performance versus the estimated outcomes? One of the things we're going to be focusing on a little is why you might not necessarily consider the same probability of default estimate in a capital measure as you might in other applications; for example, pricing.

In terms of the universal exposures, much of what you're going to hear is largely applicable to the publicly rated universe, those obligors that are rated by Moody's, S&P and Fitch. We're thinking largely about the investment portfolio where we have corporate bonds or similar instruments. The concepts are also applied to many other sources of credit exposure that you may have on or off the balance sheet as well, such as credit risk on receivables and contracts. Probably estimates could be just as easily applied to those.

I have a quick overview of those models and the main ones we're going to be talking about. The first one is the structural model, originally developed from some basic financial economic concepts in the early 1970s. Essentially, structural models are causal models that look at the probability of the affirmed asset value. The reduced form model is a very broad class of models. It's perhaps a little less specific in terms of actually identifying the cause and the reason for default as explicitly. One of the ways reduced form models are often used in actually deriving this neutral default probability from credit risk instrument prices such as bonds or credit risks. There are statistical models that really are a subset of the reduced form approach, but statistical models usually calibrated off past experience and look at predicted factors such as obligor-specific financial ratios that may also take into account economic factors calibrated on the historical data. The simplest of them all is just looking at simple historical averages as they relate to the agency ratings or the tables whereby S & P equates a AA rated obligor to a 3-basis-point default probability.

I would like to go into a little more detail on the structural default models. It starts with the very intuitive result in the early 1970s when its liabilities exceed the value of its assets. Translating that into a modeling world, one way to look at the default process might be to model that asset value out over time relative to liabilities, which are a little more fixed in nature. We would need to know the asset volatility. In reality, asset volatility is not something that's immediately transparent. You could look at the equity of the firm to get to the asset volatility. The way that was looked at was to say that equity can effectively be reviewed as a call option on the market value of assets when volatility strikes at the firm liabilities. If I think of my position as an equity holder in the firm, if the value of the assets is below those of the liabilities, in the event of a liquidation, I'm not going to get anything back. There's zero payoff on the option. As that asset value increases over the value of liabilities, I would as an equity holder theoretically have the right to liquidate the firm, sell the assets, pay off the liabilities and there would be a residual return left to me as the equity holder. The advantage of taking a modeling approach for looking at equity in an option context is you can actually use some of the standard option pricing theory to back out that asset volatility which is a non-market observable. I can look at current equity prices, the asset value of the obligor's balance sheet at this point in time and back out that asset volatility perimeter. Having done that, I could then go back to that approach of modeling the future asset value distribution.

Let's say we're looking at a point in time at some horizon in the future. We have a distribution of possible outcomes of the market value of those assets and points and market value around the book value. There is some asset growth over time. We would expect between now and one year, there will be some appreciation and natural drift in the value of those assets. The probability of fault comes from the chance that the asset value might drop below a particular default value or a market value of assets, which is such that liabilities will exceed the value of those assets. Theoretically, it is fairly simple. The trick comes in actually applying this and perimeterizing that model. There are a couple of approaches that should be taken. The first one might be to take an analytical approach so I can look at an assumed distribution around that future market value of assets. I can model it as a standard emotion, for example. I can also include a specified point of the default barrier. The original observation was for defaults when the value of assets falls below the value of liabilities, while in practice we actually see a recovery rate on debt, which implies that the value of which firm's default is somewhere above that point where assets equal liabilities. Having done that in the analytical world, it is straightforward to work through the math and come up with a solution for the probability of default from results. I'm not going to go into the detail of that. If you're interested, risk metrics credit grades is a publicly available methodology. You can go on to the risk metrics Web site, look at the paper on there and work through the map behind how they get to that result.

The other approach to take is rather than analytically specify that distribution, you can look at an empirical distribution, i.e., what we've actually observed over time. This is the approach taken by probably the most famous implementation of the structural model, which is Moody's credit monitor. The approach looked at their own proprietary base of default and effectively perimeterized the default probability by industry of their own observations about what that distribution looks like. They, I guess, made the observation that the distribution's far from normal. They were seeing defaults for six standard deviations or so away from the mean and, therefore, decided it would be far better based off of empirical results.

To give you a quick example of where that gets you in terms of actually applying a structural model in practice, the fundamental benefit of a structural model to those people who use it is that it is more risk-sensitive than for example, an agency rating and a fixed default probability associated with that. The default probability is implied from the agency rating. In one example we saw the market-based default probability reading the ultimate indication of that obligor's credit rating. In this case, they were crossing, and its rating declined every time. There is a similar sort of pattern with the equity price. As for the default probability associated with that, as the equity price trends downward, indicating that we're getting to a point where the asset value of the organization is approaching that of its liabilities, the default probabilities are trending upward. Again, it's just leading the agency rating, which tends to be a little stickier over time.

Reduced form models are something we should be fairly familiar with. Essentially, one way of looking at them is they're less prescriptive around modeling the actual default addenda as the structural models do. They rely on a mathematical process. We can effectively define the process. We have an instantaneous rate of default. One way of looking at that might be as a function of time. It can be a function of any variable. We're going to define that function. If I look at it as a function of time, a simple extension of that theory is to say, well, let's set that function of time constant for different time buckets. Default probability is something that has a term structure. As I go out further in time and look at five-year default probabilities as opposed to one-year default probabilities, they're going to be increasing. A simple application that has a rate model is to define one-year, two-year nd three-year default probabilities. You look at not just the defined case of default and no default. You can also look at rating transition effects, and the mathematics gets a little more complicated. Effectively you end up modeling this. We can look at the future state of the borrower as not just default and non-default, but we can also look at the chance that at current rating it might increase or get a better rating. An AA borrower might move to an AAA. An AA can also be downgraded to a single A or any of the intermediate states between there and jumping to default.

One of the ways in which reduced form models were first used in the way that many people think of them is actually using them to derive a risk-mutual default probability, which, again, I think is a fairly standard concept in actuarial science. We can look at the difference between spreads or prices of credit-risky and non-credit-risky instruments, for example, treasuries. That spread reflects an expectation of future default probability and recovery. We can back out what the implied default probability is from those spreads. A simple example of that would be if I were to take a risky zero two bond. I can decompose the outcomes into the chance that the bond doesn't default if I'm still a bond holder at the end of time T, and the chance of it does default with a recovery rate associated with that. Valuing those two sets of cash flow outcomes should get me to the price of the bond. If I substitute in a couple of values around that, I can settle for the remaining factor, which is the default probability. In this case it's somewhere close to 11 percent, which is I guess on the S & P scale around a single B, B- rated obligor. Theoretically, it is fairly simple and consistent with financial economic theory. For anyone who has tried it in practice, there are a lot of complications with decomposing the credit spread. The first choice you have to make is which instruments you are going to start with. I can take CDS spreads as perhaps the purest indicator of risk. I can also look at bonds. I can look at asset slot prices or any other credit-risky instrument. The complications come with the fact that there's a substantial amount of noise around the validity of those various instruments, and they may also apply to a fairly small universe as a result of that. The second complication, which we've seen a number of different academic and practical models addressed over the years, is whether the spread only reflects the expected outcomes. There are certain amounts of premium in there for the uncertainty around the outcomes as well. We don't know the probability of default as a well-defined number, so there are various multiples I use to try to address the uncertainty around that within the spread also.

Statistical models are really a subset of reduced form models, but they are expressing likelihood of default as a function of a number of obligor specific financial rations or characteristics, and also perhaps some economic indicators as well. Underlying that, the risk factors are weighted or transformed by a series of functions and weights into a single score by obligor. Having got to that score, I can either transform that into a continuous PD scale or I might associate buckets of scores with, for example, my own internal ratings and associate probabilities of fault. In terms of developing those models, this is something where the commercial banks have led the charge over the last decade or so, largely because they're less dependent on publicly observable practices such as equity prices or debt prices. They can apply to their middle-market portfolios where they have a lot more data to work with.

For perimeterization of these models, the equation itself looks a lot like a regression equation that works around perimeterizing statistical models by running a series of aggression analysis on historical data. I need a database of defaulted and non-defaulted obligors. I need to know the various financial ratios or risk factors that I'm exploring relative to those default events, so I postulate that there's leverage. A leverage ratio of an obligor is an indication of default. I would need to know what the actual leverage ratio of an obligor is if I was measuring over a one-year horizon. What would the ratio be one year before the default event? I could then look at how many obligors in a certain leverage ratio then defaulted over that horizon within my regression technique. Most of the commercial banks that have been putting these methodologies together have been using logistic regression techniques, and it's something that is very much an ongoing field of development.

It's not just the banks' internal rating scheme. Some of the vendor models that look at private firms also consider a similar methodology Moody's and S & P's credit model is similar. They look at ratios within some of your credit-fundamental categories. They're looking at predictions of fault, which might be profitability, the leverage ratios, growth or liquidity. You might also take into account some qualitative factors where you look at the banks with their large armies of credit, offices that spend a long time assessing different companies' credit. They also take into account factors such as management quality, competitive position, less easy to quantify in terms of financial ratio, but something that can be scored and incorporated in the same way as a ratio might be. The other area that some of these models try to address, which we're going to talk about that a little bit more, is the secrecality around probability of default. Probability of default is not something that is the same at all points in time. I'm not at a property casualty conference, but it's very similar to the underwriting cycle in property casualty terms.

To summarize the approaches, we've got the structural model, which looks at explicitly modeling the default event when asset value falls below a defined default barrier. With the principal model inputs, we need to know the equity price of the obligor. We also need to know information on the current balance sheet structure.

We need to know the asset value. We need to know information on the current liability structure. As a result, it only applies to firms with which we can get that information, particularly equity price. The key assumption in Kamakura's model is the way in which we decompose that observed credit spread to imply default probability. Scope obviously is for when we have observable debt prices, defaults on prices, or anything else that is an indicator of the credit spread associated with the liability structure. With statistical models, we're looking at specific indicators, as they are predictors of default likelihood. The information sources we need there are a lot more granular in terms of the financial statements of that particular obligor. Another reason why it's popular with the banks is that they have access to that information early on as that information is released. The advantage of it is when we're looking at market observables, we can look at both public and private firms. The other point to make on this is it can be used at a point in time all through the cycle measure. I am not quite so dependent on my market observables and my current expectations around evolving. Examples include implementations of that Moody's can be risk calc, S&P credit model, bank internal rating systems and George will probably touch on Kamakura's implementation.

The other point to note is that these approaches can be used in a complementary fashion. If you look at many of the banks which have internal rating systems, they may very well benchmark the default probabilities coming out of that with those out of a KMB model implied by a number of vendor models. In fact, there have recently been some industry-wide efforts to look at the relative performance of these models. There's no shortage of academic literature also looking at the performance of structural models versus statistical models over time.
I would like to talk a little more about some of the considerations in using these. With point-in-time versus through-the-cycle, it is a very important issue to understand exactly where the default probabilities are being sourced from. It's important for capital management. At the other end of the spectrum, it's also important for pricing. Think about a capital number that I don't necessarily want to change too much over time based on current expectations. The default probabilities double and the very real representation of them might be I have to hold double the amount of capital for credit risk, which is not necessarily something you want to be going back to your CEO and CFO with. You would have to tell them that this year we need double the capital because the default probability estimate has changed. It is very important to know what these default-probability tools are telling you in terms of where they're measuring and where in the cycle you are relative to where you might be in the future. I guess another consideration might be away from the capital management. The other extreme is if you are pricing. If I am looking at a 30-year contract and I want to look at some expectation of credit losses on that, I don't necessarily want to use a KMV implied default probability. Which is only looking at only the next year and, i.e., the approach might be to scale that up to 30 years to get an expected value of credit. In reality, in the 30-year time horizon you're going to go through several credit cycles, so through the cycles arrived estimates such as a historical average might be more appropriate for looking at the 30-year business cycle.

To give you an example of that, we actually apply this to, let's say, an investment portfolio. We bandage our obligors according to their agency credit rating. If I were to apply simple historical averages, I get the probability of default increasing as I move down in credit quality. I also have an application of a structural or market-based model, and for each rating band I get a number of values for different obligors within that band. In this case, they are almost all above the historical average. In some cases, they are substantially so. If I were looking at a capital estimate based on the default probabilities, it could be several times the estimate that I would get from the historical averages. Another example corresponds to a case where we have high default rate. Similarly, on the flip side, if we were in an expansion period where the default rates are low, you might see all those obligors underneath the default curve. In that situation, you end up holding less capital were you to use those default probabilities to generate your capital.

That is a brief overview of some of the default probability models out there, and some of the issues associated with interpreting those and applying them to both capital-management pricing and to other applications. I wanted to touch briefly on how we take those default probabilities and at least in the stand-alone credit risk world look at portfolio modeling. Portfolio models are looking at a distribution of outcomes. You can look at that in a couple of ways. You can either look at loss distribution where we're perhaps looking at default and non-default events in more of a book value paradigm. I simulate out my default events over the next year, I take the coverage on those and then I take a distribution of the result. The other ways to look at it is a value distribution, which is probably more applicable when you're integrating credit risk in with other risk measures. I would take some of the information that I am underwriting in terms of simulation default probabilities and relate them back into the valuation of those instruments. So as I look at credit-risky states in the future, either rating transitions or default events, that has an implication for the valuation of credit-risky instruments currently. That leads me to more of a value distribution of change in the value of my portfolio with changes in credit risk of the underlying provision.

There are a couple of approaches used for portfolio modeling. The structural approach actually lends itself pretty well to stimulating a distribution of value changes as well. In this case, the stochastic variable is the underlying asset value of the firm in question. In order to stimulate that out into the future, I need to make some assumptions around what that distribution might look like. Typically, it's very easy to assume a normal distribution around that asset value with the advantage that you can then use tail methodologies for looking at correlations between different obligors. In terms of what I need to run a structural model, some of the model inputs, you need some default transition probabilities to effectively model the uncertainty around those. Some of the commercial implementations of that can be to have the portfolio manager and risk-metrics credit manager also look at stochastic recovery rates as well on an uncorrelated basis. I also need a correlation matrix, and perhaps the most challenging part of any portfolio model is to perimeterize that correlation matrix.

The macro economic model looks at things in a slightly different way whereby you actually simulate macro economic factors out over time. Our stochastic variable is the macro economic factor that we relate to changes in default rates or changes in those factors, for example, interest rates, gross domestic product, any others. In here is the data that are similar to the statistical approach we discussed. The data needed to train that model, if you like, is that relationship over time, i.e., what do I think changes in a particular economic variable or changes in the interest rates might do to my default rate?

The closed-form model is the most analytical approach among the three. It is where I simply look at a defined distribution for credit-risk losses at the end of a particular horizon. Most commonly, people have tried to fit the default events to a distribution. The GAMMA distribution is one that seems to have come up a few times. One of the most famous implementations of the closed-form approach is CSFB's credit risk plus. If any of you are interested, again, that's an open and transparent methodology. You can go on CSFB's Web site, download the practical documentation behind there and they actually have a nice little spreadsheet example that you can play around with as well. You can put in your own perimeters, put in a bunch of your own exposures and it will come up with a credit-risk loss distribution.

Another essentially closed form-two solution is the Boswell two-credit-risk formula. Essentially, this is a structural model that's being perimeterized to get to an industry-wide application of that structural model, i.e., the perimeters are applicable and it's a one-size-fits-all approach. The model that was used in deriving the Boswell two formula was KMB's portfolio manager. They ran it for a number of different portfolio types to come up with the industry-wide perimeters they wanted to reflect in that model. Again, approaches contrast across all three of these.

This is only really touching the surface of the issue, but I would like to talk briefly on default correlations. It's not an easily observable perimeter. The way in which you might determine those correlations is dependent on the portfolio-modeling approach you might take. For the structural models, their stochastic variable is the underlying asset value, therefore they derive that correlation structure from looking at correlations between movements in the underlying asset value. A credit manager, for example, uses as a proxy for changes in asset value equity prices and that correlation matrix is derived from linking back each obligor to a weighted average of industries and the associated equity price movements for those relative to any other obligor and their mapping. They map through MSCI indexes across industries and countries. Portfolio manager takes a very similar approach. They look at industry, country and regional values. I didn't mention the bulk of these also consider a fair amount of idiosyncratic risk as well. It is very important when you're looking at the portfolio to understand that obligor defaults may not be purely influenced by these systematic factors that we've looked at. The Boswell two formula, for example, looks purely at systematic risk. It assumes you already have a perfectly diversified portfolio with regard to credit risk. That's why the regulators have looked at the capital framework around Boswell two. They've also incorporated the two considerations around banks own and economic capital models, which should have a measure or should reflect to some extent the credit risk associated with concentrations in their portfolios as well. For the macro economic model, the full correlation is to be captured by the correlation simulation of the underlying macro economic factors. I can look at simulating out interest rates for domestic products in a correlated way and then captured in my output is the inherent correlation between default events. The closed default solution, I guess it becomes less of an issue. It's embedded in the assumed distribution and the way you perimeterize that. Similar to the market-implied default probabilities, correlations are also something that you can look at on a market-implied basis. The correlations of perimeter is inherent in the valuations of CEO's basket swaps. One of the ways you can also look at correlation for different sets of obligors is to back it out of those asset prices.

We've been fairly academic in terms of going through the theory here. What it boils down to in practice is just like any other model. How you validate those estimates over time and how you ensure that your model is accurate. Boswell two has placed pressure on the internal rating system. It's also placed a lot of pressure on the vendors to prove that their models are accurate predictors over time. For those banks or entities that rely on vendor models, they have to demonstrate just as they would with an internal model how their model validates over time. Many of the vendors have been working on their own white papers formalizing those to meet regulatory standards around how their models do validate. It's still very much an evolving area. We've seen new releases of some of these vendor models over the last couple years and we shouldn't be surprised to see vendors continuing to release newly calibrated models in the future. Of course, it's still a relatively young industry in terms of the data. We have to support a lot of these models. Default is a low-frequency event. There's not always as much historical data as you would like to calibrate some of these models, although expert judgment comes into play fairly often. For those who are actually working with these models, 90 percent of the work is not around the methodology, it's around the data that you're putting into the model and you're using to perimeterize those models in the first place.

MR. GEORGE A. HOLT: I'm going to talk a little bit about default probabilities, but most of the discussion is going to concentrate more on the entire context of managing a credit-risky portfolio.

What's really involved with the analysis that's required for managing a credit-risky portfolio? Of course, we've been talking about default probabilities, and that's the first step. It's a very important step, probably the most critical in some respects related to credit risk of any of this, but there are a variety of other things that you also need to do to understand and model the credit behavior. One of those is to understand the portfolio yourself that you're talking about. We'll talk about what we mean by portfolio in a moment.

Looking at the potential recoveries in the event of default is another area. What we're really going to be getting at here is there's a whole collection of interrelated models that one needs to have to properly model and analyze a credit-risky portfolio. Recoveries on counter-party obligations are another area. Then you need to relate all of that to economic value, fair value or whatever you want to call it of the portfolio instruments if you're interested in the valuation prospective. Alternatively, our view is there are different perspectives on managing risk, which has been probably more important in the banking industry traditionally, although we see a number of large insurance companies worrying more about cash flows in financial net income as another prospective. In order to get to that, you need to understand the effect of credit defaults on cash flows in net income going forward in time.

How do you arrive at how these default events are correlated? That has an influence on the value of the portfolio. Diversification should reduce the risk of the portfolio, so you need to understand correlation of the default events themselves in order to get to that. Then you need to consider what we call credit-adjusted financial risk managers either looking at the economic value changes or the earnings changes over time.

Another typically important part is to look essentially at hedging in a credit context. That is credit mitigation, or in some cases, people want to enhance credit exposures in certain directions, and so they'll actually add the credit risk in the portfolio. When you look at an existing portfolio, you need to ask the question, "Do I want more or less credit risk associated with that and what transactions do I do to achieve that?" Those are at least some of the more important things that one has to worry about with modeling to get to the bottom line of what's the risk in expected return on a credit-risky portfolio?

Traditionally, a lot of this was done with credit-rating based models, and there are still a lot of people who are using that. The easiest way is to look at historical default probabilities by credit rating. The rating agencies publish those. You can pick those up and if you have an instrument that's rated by that rating agency as a AA, then it has a certain default probability associated with it and that's the end of the story. That is a crude way of doing it, but it's available and can be done.

A little bit more complex way of doing it is to say, "I'll think of the credit ratings as being a whole bunch of different credit states and the possible transitions that can occur between those states." I have a transition matrix with the probability of making certain transitions. One rating category is the default, so I want to know what the probability is that I'm going to end up in the default state over some period of time. That's another credit-rating-type approach. I won't say much about structural default models, which are ones that are explicitly modeling the financial structure in some fashion. That's what structure of a company is there for and you should be looking at that to explain why default occurs.

The final category is the default hazard models. There are, of course, two versions of that. One deals with historical default behavior, counter-party characteristics and economic variables that determine the future of default events. The other one is where we have market places for credit-risky instruments and we can use essentially a valuation methodology to back into an implied default probability. This is if you're familiar with option pricing to coming up with an implied volatility for an option and then using that to price other options.

I'm going to concentrate on the default hazard intensity models. They basically are economic duration or statistical survival analysis types of models that describe the probability distribution of the random time interval, so it's a time perspective on this. Many default probability models or estimates that you see are actually based on this single, fixed time period view of the default probability. Whereas these models ask a different question; that is, how much time does it take before default occurs? It's a random event that occurs. We can model that using a default hazard intensity function essentially that has some relationship to some risk factors. These models are basically one or more risk factors that could be properties of the obligors themselves, such as the financial characteristics, probability or leverage of those types of ratios that you have. They are essentially looking for things that explain in a statistical sense why defaults occur. It can be things like characteristics of the obligors such as what industry they're in or how big they are. Those things have historically been useful indicators of default. We think of the modeling approach as the counter-party economics, that is, it counter-parties default at least to some extent because of the influence of the economy. That is quite important. There are certain risk factors that probably are going to be useful explanatory variables in terms of why defaults occur, so things like interest rates, market volatility and other variables could be introduced into these models.

In terms of modeling default probabilities, ideally you'd like to attribute certain properties to them. One is that default probabilities have what you could call a term structure. That is, they differ depending upon the time period you're looking at going into the future. In my perspective, there's not a single default probability that differs depending upon how much time you're concerned with. The second thing is that the default probabilities aren't going to be the same today as they are tomorrow. They're going to evolve over time because in general the economy's going to change what those default probabilities are, and so you would like a model that can explain that. If you have that kind of model you can always get back to the model that kind of smoothes things out over the whole time period, but not vice versa. Then the third thing is that the default events are correlated over a given time period, so one needs to be able to tie together default probabilities for different obligors and understand how they're interdependent.

You don't have to do all this yourself. Particularly for public companies, you can bring up a Web page with the service that we offered that shows you the results basically of using these types of models, just purely looking at the default probability aspect of it. This is something you can go out and acquire without having to do any modeling work yourself, if that's what you prefer to do. There's a fair amount of work that goes into this.

We are using something that we've developed that's a hazard-rate model like I was describing before, with something like 13 or 14 explanatory variables in it to come up with the default probability. Another example is that you can look at evolution over time with these capabilities and see historically how the estimate of the default probability for a given period of time has developed. The model is fairly sensitive to information that's coming in on a daily basis as to what the default probability is. It's really very responsive if you like current information about default probabilities and what is perceived out there. It's interesting to compare that with the credit rating. In most of these models, one of the interesting things you'll observe is they always lead by often a significant amount. It's some sign that the default likelihood of a particular obligor is changing before the credit ratings changed. It's another reason to think about using that type of modeling approach as opposed to the credit-rating-based modeling approach.

There are also ways of computing correlations. Actually, this methodology comes up with a unique correlation between pairs of obligors for something like 16,000 firms in the database. You can get unique correlation for any pair of those 16,000. Now as I said, I normally intend this as a commercial exactly, it's more to demonstrate that this type of capability is available if you're into credit-risk modeling. You should be looking around at this or some of the competitive stuff to help you out because it's a lot of work to do it, and for public firms this can save you a lot of effort.

The next topic is credit exposures. Credit exposures are portfolio instruments with credit risk to counter-parties, and there are a whole variety of different types of instruments that one can think of and model. For example, there are the traditional debt instruments that you might find in an investment portfolio ranging from corporate bonds to even consumer loans. In many cases people use market sense derivatives like interest rate swaps or options. They use those types of instruments to hedge the market risk in positions, but the thing is they can be credit-risky counter-parties, so one needs to factor in an analysis if those are part of the portfolio we're looking at. How do those influence the overall credit risk?

More recently, although actually some of these things have been around for ages, we have credit derivatives of various types. Default swap market has become huge, but it only deals with a relatively small number of entities out there. There are a whole bunch of other things to consider, such as financial guarantees, credit guarantees, commitments and pledges. Particularly in the banking business, that's a very big way of doing business is to have pledges, etc. I think what's increasingly important are things that I call insurance derivatives, various things that enter in. For example, if you have a reinsurance contract that is essentially hedging some insurance position you have, you have that with a company that is in general credit risky. That's something that you need to think of as modeling part of this, too. There are quite a large number of different instruments and they all have unique characteristics and a lot of attributes in some cases to describe how they work. And so in each case one needs an explicit model, if you like, of how they work, how they're represented and also of how recoveries would be made in the event of default on those instruments. Recoveries can be obligor-specific, but also, in general, instrument-specific as opposed to default probabilities, which were just looking at the obligor in the modeling approach. Given that you have the recovery rates, you can figure out what the loss rate in default is.

The next step is to worry about how you come up with an economic valuation. Our approach is to use what we call arbitrage-free evaluation methods to do that. That is, you can essentially replicate the instrument with a portfolio of instruments that you understand the value of and say the values have to be the same because otherwise I can construct an arbitrage. I have an example where I'm just saying let's suppose that the lines represent the cash flows from a credit-risky coupon bond, a simple instrument, and no embedded optionality in it or anything. For that example, they can be replicated by a portfolio of credit risk. We can think of constructing a credit-risky discount rate curve that's going to allow me to place a value on each cash flow. Then I would look at the value basically corresponding to the point on that curve where each of the cash flow is occurring and I discount that cash flow at that rate. This is different from saying I have a single rate I'm going to discount all the cash flows in, because those rates can differ from one point to another and they may represent different credit risks, too.

All this is based on the derivatives valuation methodology that's been sitting around for 30 years at this point at least, and treats every instrument basically as a derivative of the underlying credit-risky zero component. If we're able to deal with the valuation of the credit-risky zero bond, then we have a way of getting that through the arbitrage-free, risk-neutral valuation of any credit-risky instrument.
To look at this in a little different perspective, I want to talk for just a second about what I call counter-party yield in spread. This is a terminology question, but counter-party yield is the yield required to hold a counter-party's instrument total, yield investor's require for that. The way we're looking at that is we're going to break that down into a default-free yield, plus a total counter-party spread. Essentially, all the credit risk is going to end up in that counter-party spread and the other part of the yield is a default-free yield. The counter-party credit spread is only in general a part of that spread, and this is what complicates some of the analysis issues in doing this, because in general instruments have different liquidity. To the extent that the instruments are less liquid, investors are going to require a higher yield to hold it. That needs to be part of the analysis if you'd like of instruments, too. Then, finally, if you're looking at a portfolio that has instruments from different taxation, you need to come up with a way of adjusting the yields to reflect them all in the same kind of basis effectively. By looking it from a yield perspective, we can convert that into a discount rate perspective, which we need for the valuation.

Looking at that counter-party credit spread, I'm going to relate this back to our default and recovery models. We're saying that spread is going to be equal to the expected loss intensity basically of the credit losses that are going to occur, and that in turn is a product of the default hazard intensity and the loss given default. We have a way of coming up with the expected rate of credit losses that are associated with a given type of instrument and a given obligor.

Credit-adjusted discount rates are obtained from the composite default-free discount rate and the expected loss intensity for each counter-party, so we can take the instantaneous default-free discount rate and the expected loss intensity and come up with a credit-adjusted discount rate. That's on an instantaneous basis. If we want to know what it is over a period of time, we have to accumulate that; integrate it over time. That gives us the kind of rate we were looking for in terms of transfer. Our risk-mutual credit-adjusted valuation approach is comprised of three steps. First, we need to determine the credit-adjusted stochastic discount factor under a particular stochastic risk factor scenario. When I'm saying risk factors here, those are risk factors that perhaps aren't totally apparent from what I've been saying. There are risk factors that influence the default probabilities. There are also one or more risk factors that influence the default-free curve in terms of the levels of interest rates at any given maturity. We're talking about all those risk factors together; the things that influence credit risk and that don't. One comment on that: This modeling approach says while the default-free interest rate is obviously telling me something about where yield curves are, the interest rates can also enter into the default probability and recovery rate models and thereby affect the credit spread. Meaning that as interest rates change, the credit spread changes also. The first step is we come up with a discount factor over a given period based on a scenario. Then we have to look at what the credit-adjusted value of a contingent claim is under that scenario. In this case, since I'm talking about a unitary zero compound bond, I just have a payment of $1 or whatever is the payment that will be made except if there's default, in which case I will get just the recovery amount. Finally, we would want to determine the expected credit-adjusted value contingent point using the risk-neutral probabilities. It sounds like people are pretty familiar with this concept, so I think it's no different than what you're used to seeing in terms of how you do this, but simply in this case using the credit defaults as part of the equation. The result that you come up with is the credit-adjusted value of that unitary cash flow, which is what we needed to do valuation of any credit-risky instrument using derivatives technology to determine it.

Given all that, one can take a portfolio of instruments and the appropriate credit risk models, that is the default models and the recovery models for each of the instruments in the portfolio, and come up with a value for the portfolio. You want to come up with a value of an estimate that also reflects optionality. If you have a callable corporate bond for example, you need to factor in the optionality in the value also with this. We like to think of it as doing valuation at the level of actual portfolio instruments, because that doesn't require you to come up with some approximation and then investigate whether the approximation's a good one in terms of the results you get. Our focus is on valuing every single instrument in a portfolio. Also, instruments are valued in different currencies in this perspective. If you have a portfolio that has corporate bonds, let's say in two, three different currencies, then you have to factor that in the analysis in terms of the value, but you want to translate it all back into what you think your base currency is at the end. The technologies available to do this can perform it for the current date, for prior dates, for future dates, as long as you have market prices and yields and estimated yield curves again. For the other risk-factor values, like foreign exchange rates and so forth, you can employ the currently available technology to do this type of valuation and come up with values that reflect the discounting of future cash flows, the value of explicit or embedded options and the expected credit losses. It's difficult putting together all the things that can influence the value moving forward.

Now I know how to put an economic value on a portfolio of credit-risky instruments, but that doesn't tell me what the risk is; that just tells me what the value is. If I want to do that, then I have to essentially simulate every evolution of the risk factors that are going to change the value of the portfolio. We have a bunch of perimeters. What this means is evolving yield curves, maybe several of them, if you have to have at least one for each currency that's involved in this. You may have more than that for various reasons. All of the risk factors describe the stochastic evolution of the yield curves, stochastic evolution of foreign exchange rates, whatever other risk factors might be out there to do Monte Carlo simulation over time of the risk factor paths. Now in reality, sometimes this isn't done in steps; it's done in a little bit different order because of efficiency concerns. That takes care of changing the environment if you like things that represent the economy or markets or whatever that influence values. The other aspect of this that one needs to worry about if you're looking at a long term prospective, meaning probably more than a month in reality, but certainly if you're looking at 10, 20 or 30 years of time is you can't assume that the portfolio that you're looking at is going to be the same. It's going to run off, for one thing. In most cases, if it's like fixed income instruments, so to make it a proper analysis, the question is what assumptions do I make about the composition of the portfolio going forward? We need a way of modeling changes in portfolio composition that deal with things like the net cash flows coming off of the portfolio and how do I re-invest those into new financial product investments, not necessarily the same ones, depending on what the strategy is. I also look at financial product investments or just new investments that we think are going to be made according to some strategy the portfolio manager has over time. Then, finally, there are usually a lot of options in portfolios these days. The marketing people really like those because it gives them a lot of flexibility in terms of features they can put into the instruments to make them appeal to the customers, but they complicate analysis quite a bit. The option contracts are going to be exercised and generate cash flows. Alternatively, they are going to deliver a new investment, that is like a bond option is going to produce at some point in the future if it's exercised, any bond position in the portfolio. All those things need to be factored in. It's a matter of having a methodology of some kind to model how the composition of the portfolio changes going forward.

Given all that, we can come up with a current adjusted portfolio value distribution by taking those risk drivers that we were talking about and simulating yield curves over time. Simulating foreign exchange rates over some period of time and simulating expected current loss intensity over time and putting that all together to come up with simulated values for portfolio instruments. That ends up giving us a frequency distribution if you like a portfolio value, both at the individual instrument level and at various levels of aggregation including the whole portfolio. What we ideally want is the whole distribution because then we can see all the aspects of it without saying I may want to know its expected value, it has to do with my expected return on the portfolio. How it's shaped tells me a lot about the risk. This approach provides you with full characterization if you want to assume you do enough Monte Carlo iterations to come up with the good frequency distribution. By the way, a similar concept can be applied to this from a financial accounting earnings perspective so that you're asking the question of what's the variability in earnings period by period going forward? If it were just an investment portfolio, market to market essentially, then you're talking about what are the gains and losses, what's that distribution going to look like in each future period going forward. The same methodology that I'm talking about can be applied to project out those distributions out into the future. Sometimes that's of interest because a lot of things are changing and you may think things look very good if you look out on a one-year perspective. But if you take it out over 10 years, maybe you will change your mind about the strategy that you have assumed. There are lots of other things you can come up with, with charts that say how much of the risk you are taking actually is credit-related risk versus interest rate-related risk. A chart can give you what we call a credit-adjusted VAR decomposed into those two aspects. Finally, that type of distribution can also help you determine what your economic risk capital requirement should look like and come up with a credit-adjusted frequency distribution for portfolio value. I'm just purely concerned with the market risk of the portfolio. How much change is that versus a distribution in the credit-adjusted value that factors in potential defaults into the analysis.

Distributions are not nice and symmetrical. They're not single modal. They often have another little blip down there where all the defaults occurred. This type of methodology gives you a great deal of power to examine what does that really look like? In other words, most of it's over here, but by the way there's this big chunk that's over here on the left side that we really need to worry about, maybe try to hedge, for example. You can also come up with risk-adjusted performance measures based on that type of distribution.

FROM THE FLOOR: Adam, you discussed reduced form models, market-implied default probability. You got a simple example of a price and in there you decomposed it so you have the recovery rate in there. How do you determine recovery rate?

MR. GIRLING: Recovery rate modeling is something that I could spend just as much time talking about as I spend on the probability of default measure, and there are a bunch of different methodologies out there. The simplest method of which is to just take historical averages as they relate to different asset causes, rating agency, security bonds versus unsecured bonds, senior versus subordinated, and effectively look up to those. More advanced methodologies, which start getting into some of the individual characteristics of instruments might look at, for example, some of the risk mitigating against those. If I have collateral underlying a particular investment instrument, I can model the evolution of that collateral over time and that may be subject to as many risk factors as the instrument price itself. I can look at how the instrument exposure changes. I can also look at how the collateral valuation might change over time. For example, if I have real estate underlying a particular investment position, that's going to be sensitive to macro economic factors as well, so I might want to model that over time as well. The answer ranges from you can go from taking a very simple approach, which just looks at historical averages, or you could get very complicated and model recovery in just the same way that you can all your others.

MR. HOLT: What we tend to do is look at what you might call unsupported positions, or where the recoveries are from, the general financial resources of the obligor, except for any financial resources, any assets that they have that might be pledged as collateral and things like that. That defines the baseline recovery rate or loss given to default rate for the instrument. For example, in Boswell two, under the more sophisticated capital approaches, there's a requirement to model the collateral on the collateral pledge. As an example, there are different types of creditors that could be used, collateral being one of the main ones, but also from financial guarantees and credit derivatives and things like that. In the case of the collateral pledges, you're going to essentially end up with a modified or adjusted loss given default rate on the instrument. But it comes from the credit protection that the collateral pledge is providing, which is related to, but not the same as, the value of the collateral. Collateral pledge is basically another type of credit derivative that has a value associated with it and it only gets exercised in the event of default. That's a very detailed and much more sophisticated way of modeling it, but may be appropriate for some situations like real estate, such as a commercial mortgage loan on a commercial property that the commercial property varies in value over time. You need to model that also. I agree there are varying degrees of this, and it depends on how important that investment is relative to the whole spectrum of things. And whether there's some requirement, for example, like there seems to be in Boswell two for actually getting down to that nitty-gritty level in modeling all those things. We're trying to provide technology that does all of that.

MR. ERIC D. HALPERN: I was hoping that one or you or both of you could talk a little bit about risk-neutral versus real-world default probabilities when it's appropriate to use them and how to interpret them.

MR. HOLT: There's a divergence of opinion on this, including the people who came up with these concepts almost to begin with, like Dr. Robert Jarrow, managing director of research at Kamakura. Some of those things end up saying they really aren't different in terms of technically or theoretically where it's appropriate to use them. I was talking about the valuation of a portfolio instrument that's credit risky. That valuation, that arbitrage-free, risk-neutral valuation process, like I said, relies on the risk-neutral probabilities. Technically, you should be using risk-neutral probabilities for arbitrage-free valuation purposes. On the other hand, the probabilities of default are going forward in time, or the simulation of the risk factors that go into that on a Monte Carlo, say multi-period analysis, those should represent real, if you like, default probabilities. As I said, I think there's a controversy in the people who do this as to whether there really is a difference between those two.