RECORD, Volume 25, No. 1*

Atlanta Spring Meeting
May 24–25, 1999

Session 35PD
Translating Bond Default Experience Studies into a Pricing Deduction

Track: Investment
Key Words: Asset Topics, Asset/Liability Management, Experience Assumptions, Investments, Market Value of Assets

Moderator: PETER D. TILLEY
Panelists: FAYE ALBERT
DAVID X. LI
Recorder: PETER D. TILLEY

Summary: Most insurance companies are not big enough to develop credible bond default experience from their own portfolio. The panel discusses methods used in recent industrywide bond default studies for public and private bonds to develop risk deductions for product pricing. Topics include:

- *Differences between public and private bond default experience*
- *Why the quality rating of the bond isn't enough information*
- *Is the risk different for bullet bonds versus sinking fund bonds?*
- *The benefits of diversification*
- *Credit swaps: Wall Street's view of default risk*

Mr. Peter D. Tilley: I head up asset/liability management function for Great-West Life in Denver. My fellow panelists are Faye Albert of Albert & Associates and David Li of RiskMetrics™. Faye has a lot of experience in the bond default world of actuarial science. She's been an FSA since 1972, and she co-authored with Irwin Vanderhoof, Aaron Tenenbein, and Ralph Verni a prize-winning paper on bond defaults back in 1988. She worked in the life insurance side of the business until about 10 years ago and now has her own private consulting firm, Albert & Associates, with an emphasis on investment consulting and the default side of investments. She will be talking about an approach to anticipating defaults of publicly rated companies.

I'll be speaking next with a presentation that looks at bond default experience studies and describes how we've implemented the experience developed in the industry at my company about a year ago. David Li is going to follow up with a Wall Street approach. He's working at RiskMetrics™, which is a subsidiary spun off from J.P. Morgan. Prior to that, David had experience working with Royal Bank of

---

Note: Tables 3–5, and the Charts for this session are not available online. Please contact Linda Blatchford at lblatchford@ soa.org or at 847/706-3564 for a hard copy.

Canada and Canadian Imperial Bank of Commerce. He has an M.B.A. in finance and an actuarial science degree from University of Waterloo. He has also taught actuarial science, most recently at University of Manitoba.

Ms. Faye Albert: I am going to talk to you about research that was done in conjunction with Zeta Corporation, an actuarial investment consultant reviewing default or trying to anticipate default experience for companies, based on information that is already known from bond defaults in the public sector. I'm going to review how we have constructed a way of anticipating defaults for a company. This doesn't cover the other step of anticipating the financial results to default because you have to take into account the salvage on any investment vehicle that defaults. My presentation covers the basic underlying probability of default, like a mortality table.

We've considered the bond weighting category that is shown by most of the rating agencies as a random set. For example, Moody's A-graded bonds is one of the categories that we considered. Each bond in that set has a Zeta score. Zeta is a way of rating companies on a totally quantitative basis. In other words, nobody from the Zeta Corporation goes into a company and talks to the management about how well the company is going to do in the future. Looking at financial results for the last three years, Zeta comes up with a numerical score. These were developed by finding scores for all the companies that corresponded to those in its data set and hierarchically arranging them. That way, any company would be able to have a Zeta rating and enter into the entire universe of these bonds. And finally to get the default rates on these companies based on the Zeta scores.

The data that we had to start with was either the Moody's or Standard & Poor's data by rating category and the Zeta scores themselves. We chose Moody's data as the base data at the time. This was in 1992. First, we chose these data because the Moody's ratings were based on amounts of bonds, as opposed to the data from individual companies, and we were interested in the probability of default on a company-by-company basis. We wanted it on a life-by-life basis as opposed to the amount of issue of debt that the company had.

Second, the people at Moody's were very helpful and willing to share with us some information on exposures that we needed in order to look at the variance.

We had the expected one-year default rate and the standard deviations for each agency rating grade. We further had that same information for cumulative 5-, 10-, and 15-year periods. What we tried to do from this information was get Zeta score intervals. We envisioned this to be a continuum of scores, so that a company could come up with any score, depending on what its results had been for the last three years. We would set the score intervals to correspond with the bond categories and then establish a mapping of these scores to particular default rates.

The Zeta scores were put into categories corresponding with the bond ratings. We already knew what the default distributions were for each of these bond ratings from Moody's. Then, in order to anticipate the cumulative default rates, we used the central Zeta scores for each bond rating; the annual default rate; and the probability of defaulting in the next year, or moving from one bond-rating category to another in our transition matrix. We played with that a little bit until it came out to be right or close to right.

Table 1 shows the first data set for the bond rating categories that we had. It starts with AAA and goes all the way down to C. It also shows the upper cut-off score and the central Zeta score for each one of these ratings. For example, rating A had a 5.5 central Zeta score, labeled as Zeta notch score. The calculation reflects the financial ratios for the companies over three years.

Corresponding to these Zeta scores, we had probabilities of a one-year rate of default (see Table 2). Let's look at BB-, which had a Zeta notch score of zero from Table 1. The probability of default in one year for that bond from Table 2 is 1.8%. That's the central point. These are the one-year probabilities of default that we got for the corresponding Zeta scores.

TABLE 1
BASED ON 1970–1995 DATA

Approx. Bond Rating	Upper Cut-off Score	Zeta Notch Score
AAA		13.0
AA +	10.85	8.7
AA	8.25	7.8
AA -	7.40	7.0
A +	6.60	6.2
A	5.85	5.5
A -	5.25	5.0
BBB +	4.50	4.0
BBB	3.60	3.2
BBB -	2.80	2.4
BB+	2.00	1.6
BB	1.35	1.1
BB -	0.55	0.0
B +	-0.35	-0.7
B	-1.15	-1.6
B -	-2.10	-2.6
CCC +	-3.10	-3.6
CCC	-4.05	-4.5
CCC -	-4.85	-5.2
CC	-5.85	-6.5
C	-8.00	-9.5

TABLE 2
PROBABILITY OF DEFAULT GIVEN
ZETA NOTCH SCORES

Zeta Notch Score	Prob. Default Within 1 Year	Zeta Notch Score	Prob. Default Within 1 Year
14	0.00010	-1.0	0.03870
13	0.00011	-2.0	0.08708
12	0.00012	-3.0	0.11320
11	0.00015	-4.0	0.13584
10	0.00019	-5.0	0.15485
9	0.00023	-6.0	0.17034
8	0.00029	-7.0	0.18312
7	0.00037	-8.0	0.20143
6	0.00046	-9.0	0.21150
5	0.00061		
4	0.00087		
3	0.00145		
2	0.00800		
1	0.01600		
0	0.01800		

Then we took those Zeta scores and did an estimation of what the five-year default rate would be (Table 3). For example, if we have a company that was a 6, then it had a probability of defaulting of 1.8%. Of the remaining companies in that group, there's a 65% chance that company would remain in that category (6) over the next year. There's a 32% probability that it would go down one, two, or three notches, and a 3% probability that it would go up a notch.

We created a distribution of what we would expect then after five years for the rating of this company to be, and the probabilities of default compared with cumulative probabilities of default for a company that was originally rated 6.

Finally, Table 4 shows the default rates based on 1970–96 data. For a BB that had a Zeta score of 0.5%, the cumulative probability of default after 5 years is 9.6%; after 10 years, it's 19.9%.

These are the results that we got, and we tried to verify them. In developing the cumulative default rates, we did the verification automatically because we went back in a circular fashion and evaluated the transition probabilities and the notch probabilities to produce the desired results, corresponding to what we saw in the Moody's data. Then, subsequently, in 1997 we reviewed data from the Zeta scores alone. We used the probabilities that we had come up with here to reproduce the results in 1997. And they were pretty good. We made a linear modification to the model to make the results come out a little bit closer, but it produced less than a 1% change. We're pretty proud of our results because they seem pretty good.

What is the usefulness of this information? It might seem that because you now have a different way of valuing bonds, what difference does it make? You could have started out with the Moody's data, so why did you need a different tool? We wanted to be able to apply this to other kinds of debt. For example, banks, when they have loans outstanding, might have companies that don't have any bonds and, therefore, have never been rated by Moody's. We can use the Zeta approach to assess a numerical valuation to these companies. We can apply these default rates and get some idea as to what the financial experience will be. It can also be expanded to evaluate the riskiness of a portfolio of obligations of different classifications.

Mr. Tilley: How can we take what Faye has just explained and turn it into something that we can use as a pricing deduction? Without getting into the specifics of pricing deductions for my company, or any of your companies, how can we take these kinds of experience studies and turn them into the practicalities of saying, "It's *X* basis points for a BBB bond with a seven-year average life that's senior secured debt?" If we don't take those kinds of pricing deductions and pass along the full coupon of the bond to our pricers, sooner or later we're going to find ourselves short. After listening to the keynote speaker today, I wondered if anyone was going to show up here, because the speaker predicts that there aren't going to be any bond defaults. Everything is going to be so fabulous, unless we're investing in Japan or countries outside the U.S. But the reality of it is, even in a good economic environment, like we're in today, there will be defaults.

To get a default loss, we have to look at it in two separate stages, and Faye gave one example of coming up with some very detailed analysis on how we can look at the first step, which is to estimate the probability of default. For a particular bond that you have, what's the likelihood that it is going to default this year, next year, and the year after that? But when it does default, there's a second piece of information that we need: an estimate of the loss severity if default occurs. When I was a young actuarial student, I figured if a bond defaulted, that was that. You lost all your money. But that's really not the case. There's a salvage value associated with these. And Faye alluded to that as well.

Depending on the markets at the time that this bond goes into default, and the situation that the company is in, there's a value to that bond. It might only be 50 cents on the dollar, 40 cents on the dollar, or 70 cents on the dollar, but you haven't lost all of your money. Losing all of your money is the absolute worst-case scenario.

What kind of experience studies are out there to give us some clues as to what to assume for the first part of this calculation, the probability of default? There are studies that you might have seen from Standard & Poor's, which does an annual study that looks at the experience on publicly rated bonds. Moody's also does the same sort of thing on an annual basis. You'd think that they would come up with identical results but, since Moody's has slightly different rating categories and different ways of looking at things relative to S&P, you actually should look at both studies.

SOA does a private placement bond experience study that is more focused; it's private placements only, as the name would imply. I think it started out as a private placement and commercial mortgage study, but I suspect the SOA found that commercial mortgages were too heterogeneous across companies, so you couldn't do a really good comparison. And, there was a loss recovery study written by Edward I. Altman and Vellore M. Kishore. Altman is a professor at, I believe, New York University, who has done extensive work on bonds that are below investment grade.

The S&P study covers experience from 1981, and it develops the transition matrix method that Faye showed us on the Zeta scores. You look at where a bond is now and determine the probability that it is going to move to a different rating category by the end of the year. If a bond is now a single A, there's a lot of inertia. The most likely event is that, by the end of the year, it will still be a single A. But there are possibilities that it could be upgraded to an A+ or AA-, or it could be downgraded to a single A- or a BBB+. It covers the whole gamut from AAA down through D for default.

There's also a not-rated (NR) category that covers issues that are either paid off or the ratings were withdrawn for a variety of reasons. A bond may be a single A bond at the beginning of the year, and an NR at the end. We have to have a way of adjusting for those situations because we don't want those in our sample any more.

The Moody's study takes a different approach. It goes all the way back to 1920, so there are some pretty interesting statistics from the late 1920s–30s in there. It develops one-year probabilities of default rather than a transition matrix, and determines, if a bond is rated single A now, what the probability is that it's going to default between now and the end of the year. Instead of looking at the gradations of being downgraded or upgraded, you can develop select and ultimate default experience from the time of issue. It looks more like the mortality tables that we're used to dealing with. If you take the analogy of bond defaults and mortality a little further, the S&P study looks at people getting sicker and sicker and dying, whereas the Moody's study says they're either living or dead; it's a binary situation.

SOA's private placement study currently only covers experience from 1986–94. The data to update through 1998 is now being collected. It's a painstaking process. I think it's going to be at the least the end of the year before the committee members finish collecting the data because they ask for the data in many different ways, in a fashion that allows them to do several interesting calculations, as you'd expect from an actuarial group.

This study covers about 40% of the private placement holdings by life insurance companies. It's not as comprehensive as SOA mortality studies, but it still gives us a good grounding on the private placements that are held by our industry. One of the conclusions that I took from reading this study, and this is more perhaps my opinion then SOA's specific conclusion, is that default losses on investment-grade public and private bonds are similar. In other words, I can substitute private placement studies for the public studies that are done by S&P and Moody's.

The Altman and Kishore study was published in the November/December 1996 issue of the *Financial Analysts Journal*. The article was titled "Almost Everything You Wanted to Know About Recoveries on Defaulted Bonds." It looked at about 700 default events from the period 1971 to 1995. When you're only looking at defaults instead of a universe of bonds, you have to cover an awful lot of territory to get a statistically credible number of events. It's kind of like looking only at deaths, instead of an entire population.

The main point that I wanted to make to you from this study is that the only significant factor to determine the recovery rates and the salvage values on bonds is seniority. Once a bond has defaulted, what is it going to be worth at that point? The only significant factor that they found was the seniority of the bond. It didn't matter whether it was originally rated AAA, and it didn't matter whether it was a seven-year bond versus a three-year bond when it was first issued. What mattered on that salvage value was whether it was senior secured debt, senior unsecured debt, secured subordinated debt, or subordinated debt. There were a few other specialized categories in there, too, such as equipment trust certificates. But I think senior secured and senior unsecured would be the ones with which we're most familiar, given the portfolios that our companies have.

How can I take this information and turn it into a default deduction for new bonds that I'm acquiring? If my investment department phoned me yesterday and said, "We bought a seven-year average life private placement bond; it's senior secured, and it's got a coupon of 6.5%," I then have to take that coupon and translate that into a yield that I'll communicate to my pricing folks. What deduction do I take for bond defaults?

And now I'll couch my presentation in terms of suggestions. This is, as you can probably guess, the way my company does things, but it's certainly not the only possible way. You could use a Standard & Poor's-style transition matrix. Rather than looking at that one-year probability of default, you'd look at the probabilities of moving from a particular rating category to a different rating category by the end of the year. You could use recovery rates from Standard & Poor's, Moody's, Altman and Kishore, or even some estimates that your asset portfolio managers would have for you. Perhaps their recovery rate assumptions would be a little higher. What we found in looking at these various studies is that there was a lot of similarity in the recovery rates and the salvage values that were reported in these various data sources. We ended up using one of the original three. We didn't go with our asset managers' estimates.

With that kind of information then you can run thousands and thousands and thousands of Monte Carlo trials, as long as you're assuming the independence of default experience. (And I'll get to an issue with that later.) For each trial, start out with a number of generic bonds—for example, 100 bonds with a variety of par amounts—to try to simulate your investment program. If you have 100 bonds, they're not all going to have $1 million par value. If your typical investment program is to invest in some $3 million, $5 million, and $2 million pars, put that kind of mix in and see what that does to the distribution of overall results. With the projections, capture a pattern of annual cash flows for each trial. As you're going through this trial, each year you're going into that transition matrix and saying, "What is the current rating on this bond?" Maybe it started out as a AA and through this particular Monte Carlo simulation it stayed a AA for the first three years, then it was downgraded to a single A. And, if it's a long enough bond, maybe it's actually worked its way all the way down to BB, a single B, a C and then, all of a sudden, a default.

But, even if it started out as AAA, and it's not that long a bond, if you do enough Monte Carlo trials there is still some small probability that that bond is going to default within the next two or three years. That's what that transition matrix will tell you—that there's a lot of inertia around the current rating, but it's certainly possible to go into default, even from a single A or BBB category in a short period of time. If the Monte Carlo simulation says the wrong random number came up and this bond defaulted, now you put a salvage value in for that, and that becomes the cash flow for that bond. You can put in the coupons that you expected up until the time that it defaulted and the salvage value when it does default.

Look at the internal rate of return of the cash flows and compare that to the starting bond yield to determine the default loss for that trial. Maybe you had a bunch of bonds that started out with 6.5% coupons, but when you looked at the pattern of the overall cash flows, it came in at 6.42%. I conclude that the difference is eight basis points loss to default. That's the only decrement that's going on in my model. Develop a distribution of the default losses for this set of trials and run separate sets of trials for each term. In other words, run them for 5-year bonds, 7-year bonds, and 10-year bonds. Also, run separate trials for each initial quality rating—AAA, A, BBB—and for each seniority. The seniority comes into play when you're looking at the salvage value.

What kind of sensitivity tests should you do? Studies for the salvage value show that those recovery rates can have a fairly large standard deviation. Again, even Altman and Kishore, who looked at 700 default events, found that all 700 weren't in one particular legal classification. They weren't all senior-secured; some were senior-unsecured. And you don't necessarily have a really tight distribution of salvage value. You should test with a standard deviation that they give you in these studies, and you can stochastically generate these salvage values. To do that, you need a lot more trials.

At my company, we multiplied the number of trials several times over to get what we thought were credible data, once we started putting in a standard deviation on the salvage value. That is because you don't get that many defaults to begin with and, if you want to put a standard deviation on top of that, you need to have that many more trials to get some kind of a distribution. We found that, although that loss distribution was wider, it wasn't drastically wider. I don't know whether I can conclude that this is a guaranteed true statement, that we just didn't do enough trials, or that we happened to pick just the right lucky random numbers.

Another sensitivity test that we did was to test different patterns of cash flows. My company invests in lots of different kinds of bonds. Some of the bonds are bullet bonds, so a five-year bond pays coupons for five years and then the principal comes back all at once. We have other bonds that are considered five-year bonds because they have a five-year average life of principal. Notice I'm not talking about duration now. If I have, let's say, equal annual payments between years three and seven, the average life of principal will come out to a five-year average life. This is a sinking fund bond. For that same average life, we found in our testing very little difference in the default risk between sinking fund bonds and bullet bonds. We attribute this to offsetting events occurring. Although the sinking fund bonds have longer maturity dates, they give you a higher default risk—a higher chance that this bond ultimately will default the longer it goes. The fact that it's paying down the principal in the meantime gives you a lower exposure in those later years. A combination of a lower remaining principal balance and higher probability of default seems to balance out.

That's what we did for new asset acquisitions. Now let's look at your existing bond portfolio and getting some comfort level about the right provisioning for bond defaults on the portfolio. Do a similar analysis using Monte Carlo trials, look at your entire portfolio of bonds, and use the remaining term of each bond. If it's a seven-year bullet bond that you bought two years ago, it now looks like a five-year bond, as far as your default risk is concerned. Use the current quality rating, of course, not the original one.

And here was the tricky part for us: Use the seniority class, although it may not be available on asset administration systems. You might change that word "may" to "will" because it's not the sort of thing that we generally kept track of from an asset management or an A/L standpoint. However, once we got into this project, we found we needed that information, and fortunately we're not as big as MetLife or Prudential, so it only took one person a few months to go back and research each bond. So look back into the paperwork on each particular bond. Figure out whether the bond is senior-secured or unsecured.

You can compare the loss distribution you get with your current level of deductions. If your administrative system looks like ours, there's a field in there that shows what's being taken off for defaults each year. You can compare either year by year, or on some present value basis, whether the deductions that you're building in are matching up well with the distribution of expected losses that you have from your Monte Carlo testing. And of course, you'll have different standards, depending on whether you're doing this for reserve adequacy or solvency.

It's just like any cash-flow testing exercise. If you're doing it for reserve adequacy, you want a pretty good comfort level. If you're doing it for solvency you want an excellent coverage level on your distribution. And you might decide that you need to do some extra analysis on the left tail of the distribution. Some companies might find that their bond portfolio could use a little more diversification. If you have concentration in a handful of bonds with some fairly high par values, you might find that that's giving you a bit more of a left tail then you'd like to have.

My final comments lead in to David Li's presentation. Is it really safe to assume the independence of bond defaults? That's what my presentation so far has assumed. If we don't assume independence of bond defaults, we get very much more interesting left tails. But what kind of data can we use to get these kinds of dependence or correlation factors? Should we or can we even model defaults that depend on the prevailing interest or the economic environment? We've looked at default studies that take us back as far as 1920. The SOA study was more recent, but went through some interesting times in the earlier 1990s, where private placement defaults were a little higher. What kind of time period do you need to look for and what kind of left tail coverage do you need to make sure that you can handle it? It's almost like a reinsurance question. If you look at the amount of provisioning that you're taking and 1930s kind of default experience, are you still okay? What multiple of recent experience are you covered for?

Finally, can we use test results to analyze the effects of diversification rules in our companies' investment policies? My company and probably every other company out there have rules about how big a particular loan obligation can be, relative to the size of the company and to other parts of that sector or industry that we're lending to. We might say that on a single A bond we won't lend more than $*x* million dollars, or we don't want more than $*y* million of exposure in this particular sector. And sometimes that costs us money. There's an opportunity cost if a particularly good bond is violating some of our diversification rules. I think this kind of testing can be a valuable tool for looking at those diversification rules and determining just what kind of risks we're taking, whether we're leaving any money on the table, or, heaven forbid, whether we've gone in the other direction and maybe aren't diversified enough.

Mr. David X. Li: I'm going to talk about RiskMetric's™ approach to credit portfolio management. Our firm is a spin off company from J.P. Morgan. Five or six years ago the ex-chairman of J.P. Morgan issued an order to this group, essentially saying "Send me a report every day by 4:15 p.m. outlining the bank's position and a possible worst loss in the next three days." The group leader came out with a methodology called "value at risk," And that concept has been popularized by our technical document called "RiskMetrics™." The group is J.P. Morgan's external risk management service and products. We get a lot of inquiries from other companies for a document that was issued and adopted by the Bank of International Settlement. And we provide support for them to use our concepts and software. The group decided to spin off from J.P. Morgan last year, so I'm now going to talk about RiskMetrics™.

Currently, in the banking industry, the focus has been shifting from market risk management to credit risk management. I will present our group's approach to credit portfolio management. The methodology is called CreditMetrics™, and it contains three components. First, it's a methodology. We have a technical document issued two years ago, which anyone can download from our Web site. Second, we provide the data. When it comes to credit management, data becomes very precious. You can get intense data and you can get tick-by-tick data of the stock of IBM. But when it comes to credit, the data is pretty scarce, and the quality of the data is such that you have to have some doubt about it. Third, we have software called CreditManager. CreditMetrics™ essentially provides a framework for managing a credit portfolio.

What does CreditMetrics™ do? If you are a vice president of an insurance company, it's very likely that you manage a large portfolio of fixed income assets. You need to ask yourself a lot of questions. First, you want to identify concentration risk. How exposed are you to specific credit events? Next, you need to allocate capital or reserves to different lines of business or try to put some money aside to prepare for future losses. How much should you carry in reserves? And sometimes you are wondering about the marginal cost of specific transactions. What is the marginal capital cost?

Also, you need to put in an optimization rule. How do you get a maximum return at a given level of risk? For the banking industry, you need to set up a limit, so you cannot let one trade go wrong and bring down the bank. You need to ask the president how much exposure is too much for the whole bank or for a specific trading unit.

And you also need to do some sensitivity analysis. Suppose I get a portion of a specific bond. How much impact does it have on my total portfolio? Finally, the latest development in derivatives is called credit derivatives. People try to use a market approach to rearrange their credit risk. And a lot of insurance companies are either investors or counterparties of those transactions.

Who develops CreditMetrics™? It was developed within J.P. Morgan, but has 25 cosponsors from different sides of the business. The cosponsors can be placed in four categories. The first group is the largest banks. It is in a banker's interest to manage his or her trading, loan, or bond portfolio actively. We have Bank of America, UBS, and the three largest Canadian banks: Royal Bank of Canada, CIBC, and Bank of Montreal.

The second group is non-bank institutions from the trading side, or loan portfolio management, and they look at it as an asset cost they can trade.

The third group of sponsors is rating agencies. All the major rating agencies are cosponsors. The fourth group includes some of our competitors, like KMV or C*ATS, which essentially do exactly what we are doing—provide risk management solutions to financial industries. When it comes to credit, there's a consensus that people need to work together to solve this problem. CreditMetrics™ is committed to building the benchmark for credit-risk measurement.

CreditMetrics™ was launched in April 1997. We've printed 20,000 technical documents. At that time, we had five cosponsors, three data providers—including Moody's, Standard & Poor's, and Reuters—a CreditMetrics™ Web site, and a prototype version of CreditManager running in Excel. But we had no users.

Today, we only have 3,000 technical documents left, 25 cosponsors and 7 data providers. Every week our Web site has 1,000 inquiries. Our software CreditManager is in version 2.5, and we're going to issue another version soon. And we have 75 CreditManager users. Those are people who actually manage their bond, loan, or fixed-income portfolio in banks, insurance companies, and some mutual funds, too.

Currently, CreditManager can take the following instruments: bonds, loans, receivables, letters of credit, and loan commitments. We also can cover market-driven instruments. We call them monthly debit industrials, which include derivatives like swaps and forwards. We can also cover the basic credit derivatives like credit default swaps and total return swaps.

What kind of information do we need in order to start a credit portfolio? For example, for bonds, first you need to know the current rating. Then you need to know the credit spread, because, when you talk about the difference in rating at the end, that will reflect on the pricing. The bond's price translates into yield to maturity. Comparing the bond's yield to maturity to the corresponding Treasury's yield to maturity gives you a spread to the Treasury. That's the information you need. When a AAA bond gets downgraded, the spread will be higher and the price will become smaller. And it's similar for some other instruments.

Then you try to run a Monte Carlo simulation. You are given a transition matrix, which is probably taken from S&P historical studies. The matrix in Table 5 has seven credit states plus one default state. If you read across each row, it gives you the transition probability. For example, a AAA bond would have a 91.35% probability of staying AAA after one year. The default probability for AAA in one year is very small. It's about one basis point. This transition matrix gives you the probability of a bond moving from one state to another state.

Next, we study the change of each instrument. Suppose you are holding a BBB bond today with a one-year horizon (Chart 1). At the end of one year, this BBB bond can have eight possible states. It can get upgraded to A, AA, or AAA; remain the same; get downgraded to double BB, single B or CCC; or it can even default. For each of those scenarios, we calculate the instrument's value. At the end of one year, we would price the bond using corresponding spreads to do the discounting. That's how we get the market value of a bond at the end of one year. Suppose this bond today is a BBB. If it gets upgraded to AAA, the spread will become smaller and the price would become higher. That's why here you see, after one year, the price becomes 100.9% instead of 100%. If the bond is downgraded, the price would become smaller, and in the default case, we would get the recovery.

As Peter mentioned, there are a lot of uncertainties with the recovery or salvage value. In the U.S., for AA rated bonds, the mean recovery rate is about 40%. If you look at the standard deviation, it's also about 40%. When we try to build a loss distribution for that, we have to consider the uncertainty associated with recovery. What we do is model the recovery rate as a beta distribution and match the mean and the standard deviation. As actuaries, you will all understand why we use the beta distribution. The beta distribution gives us a value between zero and one, which can be used to describe the recovery rate.

How do we study correlation? As actuaries, in our training for insurance risk, we can assume the risks are independent in a lot of cases. The survival or death of my life has nothing to do with anyone else, so that's a pretty fair assumption. But when it comes to default, we can reject this assumption right away. If you look at the historical default rates, defaults tend to cluster. One year of high defaults will usually be followed by another year with high default probabilities. Obviously you see some kind of correlation there. The question is, how do we study this correlation?

Now we'll take one step back. We feel the default correlation comes from fundamentals. Our philosophy is that the default correlation comes from asset correlation.

How do we study asset correlation? We use stock price to approximate the asset value of a company. And, the correlation of default can be translated into the rate of return correlation. That's the approach that we take. Essentially we map the default into a normal distribution. For example, we might have two obligors. We translate them into a normal distribution. If they tend to map together, they're highly correlated. If they are far away, they are less correlated.

We map each company into certain industry, sector, and country indices. We study the correlation between the asset correlation then translate it into default correlation. There's a mapping from asset correlation to default correlation.

Next, we put the market data and the portfolio of exposures into the CreditManager software and run a Monte Carlo simulation. At the end, you can do all the analysis we required at the beginning to try to answer all those questions. How do you allocate capital and reserves? How do you identify concentration risk? How do you study marginal costs? Do you set up a limit for specific portfolios or try to optimize your whole portfolio based on return.

Let me just show you some examples of how we do that. First, since we have calculated the loss distribution or the portfolio distribution at the end of a horizon, and we did the whole distribution, then we can get any information we want out of it. The first question relates to a value-at-risk idea. Under 1% of the scenarios, what's the worst loss you could have within one year? With this distribution, you need to get the break point where the tail part will give you 1%. We can also do some concentration analysis. Suppose you are a manager of international bonds, not just U.S. bonds, in many sectors. We can do a marginal plot such as the one in Chart 2. If you look at this chart, you can certainly identify that the majority of the risk comes from the U.K. telecommunications sector.

You can also do some return-and-risk analysis. The plot in Chart 3 gives you the absolute exposure size and the marginal standard deviation. Here you can identify those with a small size and high risk. But which are those with a high size and low risk? Which are those with a really large size and a large risk? You can identify each one of them and also do some return analysis. Chart 4 translates the graph into return. You can know exactly which securities are on the high-risk, low-return side.

At the end, you can also do some credit derivative pricing. The bottom of Chart 5 gives you the total portfolio distribution at the end of one year. You might say that this is a pretty good portfolio. In a sense, the portfolio value concentrates on the mean. But still, as in a majority of all credit portfolios, it has a small probability to have a large loss.

After you look at this distribution, suppose you are a portfolio manager and decide to buy some protection to protect it from extreme loss. You have a payoff function (top of Chart 5) that says I don't want to cover all possible losses, I want to cover only a loss above level D. Using our software, you can value how much premium you should pay for one year of protection.

Chart 6 shows the pricing as a function of percentile loss you try to protect against. If you try to protect against more loss, you have to pay a higher premium.

Mr. John S. Bath: Peter, in your studies, you translated your default experience into a cost. Did you evaluate at all a capital cost associated with perhaps higher risk-based capital (RBC) factors as a bond gets downgraded? Or is that just viewed to be rather incidental and not part of the calculation?

Mr. Tilley: I have seen studies that looked at relative RBC costs, and the kind of return that you need on capital to support the extra spread you would get on different kinds of bonds. Because you have to hold a higher capital factor on a BBB bond than you do on a AAA bond if you want the same kind of return on your shareholder capital, you should get a better return on that BBB bond than you do in the AAA. The work that I have seen seemed to show that there was very little relationship between the RBC factors and the extra spread that you got on the market. Even though spreads are wider now than they were a year ago, you just don't get the extra spread on a BBB bond that the extra capital requires. And that's a tough decision for a company. It has to decide whether its going to stick to investing based on a rather arbitrary set of NAIC factors, or conclude that it has lots and lots of capital, so defaults really aren't that big a concern and the RBC ratio is already pretty good. It could be a little better if they all went to AA and AAA bonds. But there wasn't a good mapping between the extra spread that you got as you went down the credit curve, and the extra return on capital that you were suppose to get with that spread.

Mr. Bath: Peter, I'm not sure if your answer was exactly related to my question or not. Let me just make sure. Say that you have a BBB credit exposure and you monitor the experience with the upgrades and downgrades and so forth. Eventually, it defaults or doesn't default and there's a recovery. You can quantify the present value of that actual default cost over the experience period that you're monitoring. But what I'm wondering is, let's say it's a single A that gets downgraded to a BBB in the first year, but pays off nine years later and there hasn't been any credit loss. There still could be considered to be a capital loss associated with the fact that, at the beginning of the second year, you had to hold more capital against that BBB. I was wondering if you evaluated that at all in your studies?

Mr. Tilley: Based on what I've seen in terms of these RBC factors versus the spreads you get, we determined that it's just not worth it to drive our investment policy by RBC. Now, having said that, we also don't go into asset classes that have 30% RBC factors. Common stock would be a good example of that. If companies are actively in that below-investment-grade bond market, I think the capital then becomes a much more important consideration for them. My answer refers more to the AAA to the BBB- investment grade category. The different factors don't seem to relate well to the different spreads in the market. And we have essentially decided to ignore the factors as part of our capital management process.

Mr. Scott D. Houghton: My question is related to the last one. Have you considered a case where a bond is investment grade when it's purchased and then downgraded to a point where it's not investment grade anymore and the insurance company, for whatever reason, decides to sell it? Maybe it doesn't want to hold noninvestment grade bonds and is basically taking loss at that point. Have any of your (or the other panelists') studies considered credit losses like that? And what can a pricing actuary can do to reflect those types of credit loss or credit risk events?

Mr. Tilley: We've considered it, but we haven't implemented it. It really depends on your company's strategy. If your company feels that a bond being downgraded, as long as it hasn't defaulted, is not a major problem, and if it's taking a buy and hold strategy, then you wouldn't worry as much about that. If, on the other hand, you're looking at total return and actively managing your portfolio and doing lots of trading in the portfolio, then I think you have to take that kind of thing into consideration.

What you would do with the techniques that I described is, instead of waiting for the default, you'd have some other wall that you hit. If the bond is downgraded to below investment grade, you'd have an assessment of what the market value of that bond would be—the analogy being the salvage value if it defaults. And you'd change your modeling that way. You would probably see a lot more frequent credit events with lessened severity.

The data that SOA collects as part of its private placement study tries to address these kinds of issues as well. You're starting to get a flavor now for how complicated this data would be. You need to look at all kinds of different credit events, not just defaults, but this somewhat steady erosion of a bond on its way down and what's happened to the cash flows on the bond as a result.

Ms. Albert: Going back to the low-tech approach that I have often taken, I think that many companies try to incorporate a charge into their pricing, like Peter was talking about. It's a small charge against the anticipated return on the bond to account for that kind of loss.

Mr. Tilley: Another sort of loss occurs if a bond defaults at a time when interest rates are quite low. You do get a salvage value, but you have to reinvest the salvage value in a lower interest environment. Maybe you've gone from $100 of par at an 8% coupon, to $60 of par at a 4% coupon, which is a double whammy. Unless you can build in a correlation between the probability of default and the economic environment in which it all happens, it's going to be a very complicated modeling exercise.

Mr. Li: It seems to me there are two perspectives, either from investment banking or trading floor points of view. Every morning, when you come to the desk, you know what the portfolio value is that you have. When there's new information about this bond, sometimes even if there's no downgrade or upgrade, there's still some change in the spread. That will affect your present value. For an insurance company or other company that takes a buy and hold strategy, whether a bond is AAA or single A, makes no difference. The bond still pays a coupon and at maturity you get the principal, so it doesn't have an impact on you because you don't try to mark everything to market. I guess that's a difference.

The next question is, how about a capital charge for RBC purposes or capital allocations to cover defaults? It seems to me that more and more people are trying to move from this accounting approach to a market approach.

From the Floor: I have a question for David Li. In your presentation, you had a BBB bond that you split off into eight possible states afterward. And you assigned a probability to each of these states and an instrument value for each of these probabilities. I noted that if I were to take the expected value of the instrument after those probabilities, it's less than 100%. Therefore, this model would indicate to me that the expected value of my bond is less than its current value. I don't know if there's any concern about calibration when it's modeled such that it would give you a different value for the bonds from its current market value?

Mr. Li: As you mentioned, this number is mainly for illustrative purposes. For example, this instrument's value is at end of the horizon. If it's a one-year horizon, we don't even give a recovery value there. Even though it gives a recovery value, the bottom part is a value at the end of one year. It's still possible that if a bond was par today, you could have two payments of a coupon, and by the end of one year the bond could be worth less than $100. The approach we are taking is the market approach. At the end of one year, by using this simulation, you know what the rating group this bond is in. Corresponding to this rating group is a spread. We use that spread curve to do the discounting of all the future cash flows.

From the Floor: Does it still mean that, even if I take the coupons at the end of the year, I'll get a coupon in between and use this new interest rate to calculate my value? But, say I want to know now, using this model, the value of this bond? Should I use these values after one year and basically discount them back one year to get the value of this bond? And, is this something that you need to do—calibrate your model using these default rates so that the value of your bonds that you get right now is equal to the price of the bond? Otherwise, your model would tell you not to buy this bond.

Mr. Li: That's an interesting question. Essentially there are two ways to calculate it. We call it a risk-free forward bond price. Given the spread and the yield curve today, you can get the risk-free forward bond price. You can do it exactly as if the bond were risk-free. Alternatively, you can see the eight states and the capital value at each state. Then you calculate the expected value. Your question is, do these two approaches give you same value? They don't, in general. In a lot of cases, if you want to do a calibration, you calibrate the transition matrix. You make adjustments to the transition matrix to add a constant to it, so that the two will give you the same value. When you risk-neutralize the transition matrix, the two will give you the same value. For the larger portfolios, we don't do this calibration. But when it comes to pricing some specific structures, like a credit default swap, that's a first requirement we have to satisfy.