RECORD, Volume 29, No. 1*

Washington, D.C. Spring Meeting
May 29–30, 2003

Session 19PD
Credit Risk Modeling for Life Insurers

Track: Investment

Moderator: MARTIN K. LE ROUX
Panelists: WALID SHINNAWI†
MARTIN K. LE ROUX
LARRY H. RUBIN

Summary: Over the last decade, a number of large financial institutions have implemented sophisticated models to quantify, aggregate and manage their credit risk exposure. These models are also used for active portfolio management and capital allocation purposes. To date, most users of these models have been banks, but life insurers can use the same models. The concepts and techniques underlying these models have many parallels to actuarial work in other fields. Presentations cover the differences between quantitative models and traditional techniques for measuring credit exposure; pricing models versus portfolio risk management models; and an introduction to some commercially available models, such as J. P. Morgan's Credit Metrics and Moody's/KMV Portfolio Manager.

MR. MARTIN K. LE ROUX: This is session 19, "Credit Risk Modeling for Life Insurers." I'm Martin le Roux with ING Institutional Markets in Denver, Colorado. With me on the panel today are Larry Rubin, a managing director with Bear Stearns in New York City where he provides advice to life insurers on balance sheet management, and Walid Shinnawi, a director in the client solutions global group at Moody's KMV in San Francisco. I'm going to start by giving some general remarks, and then I'm going to hand it over to Larry and then Walid, who are going to talk in some more detail about this topic. What you see on the screen is a little bit different from the material that was put on the SOA Web site. So, if anyone's interested in a copy of the presentation that you see on the screen, just send me an e-mail, and I'd be pleased to forward it to you.

---

†Mr. Walid Shinnawi, not a member of the sponsoring organizations, is a Director at Moody's KMV in San Francisco.

I think most people here will be aware of the importance of credit risk management to life insurance companies. Chart 1 shows some recent statistics from Standard & Poor's of investment grade defaults. More specifically, they show the percentage of corporate bonds that were rated investment grade as of January 1st each year and that defaulted in the following 12 months, going back some 20 years.

There are a couple of things that stand out here. One is that, on average, investment grade defaults are quite low, probably less than 10 basis points, but there's a huge amount of volatility, as you can see from 2002, and so there's risk here that you can't diversify away simply by adding more credits to your portfolio. It's this risk that credit risk models are really trying to get at. This is also from S&P. Chart 2 shows what happens when you add downgrades as well as defaults into the picture. It's a very similar picture qualitatively.

Credit risk models have been developed to address a number of issues around credit risk. One is simply trying to forecast the level of expected losses that you might see around a portfolio. I'm not going to say too much about that today. They also attempt to quantify the risk of unexpected losses, losses in excess of the mean, and try to quantify the amount of credit risk capital that should be held around that. They do so in a portfolio context, meaning that instead of just looking at each credit in isolation, the idea is to try and capture dependencies between the different credits that might make up a portfolio. It's really these last two points that I'm going to focus on.

Before I say more about credit risk models, I want to contrast these models with traditional approaches to managing credit risk capital. By that I mean approaches such as NAIC risk-based capital, the Canadian minimum continuing capital and surplus requirements system and the S&P life company model. These approaches have in common that they're all factor based. They tend to be driven primarily by credit ratings, and this raises quite a few issues. One is that it's often not cared exactly what these numbers represent, what level of confidence they're set at, or how they were derived.

The factors tend to be independent of maturity dates. A 1-year asset would be considered no less risky and no more risky than a 30-year asset, for example. Most importantly, they look at each asset in isolation. In general, these approaches don't capture how well or how poorly your portfolio might be diversified. The factors stay the same over the time. They don't respond to current economic conditions. It's a very broad-brush approach that all BBB assets, for example, would be considered to be the same, although in reality a single BBB low can be a lot riskier than a BBB high, for example.

It's in response to some of these issues that a number of credit risk models have been developed, and I'm focusing here on portfolio models. There are quite a few other models out there that address other aspects of credit risk. Probably the three best-known ones are CreditMetrics, KMV's Portfolio Manager and CreditRisk+. There are others out there. These three are the ones most frequently encountered. The details in the models are very different, but they have a number of elements in common. They're trying to model the portfolio loss distribution, and they're trying to capture the idiosyncrasies, the particular portfolio that you're looking at, as well as systematic risk, the component of credit risk that's going to be present no matter how well your portfolio is diversified.

Chart 3 is a generic picture of a portfolio loss distribution. This might be produced by any one of these models. As you might expect, it's a very skewed distribution. There's a high probability of relatively low losses and a low probability of very high losses. The portfolio standard deviation, the loss distribution standard deviation, is often referred to as the unexpected loss, and most institutions using these models typically set capital by reference to some high-end tail-risk measure, some percentile measure out in the tail, and it's usual to set capital by reference to the difference between that percentile and the mean of the distribution. Institutions don't typically hold capital against expected losses. The view is that that should already be covered off by margins in your spread.

There are a couple of things that you have to think about when specifying a loss distribution. One is the time horizon over which losses are modeled. One year is a very common, almost universal choice. The view seems to be that anything shorter than that is not really going to be meaningful, and anything longer you run into very severe problems in trying to parameterize your model. The question is what confidence limit is appropriate? The common view seems to be that this ought to be consistent with the desired credit rating for the institution. If you're an AA institution, historically AA institutions have about a five-basis point default rate over a 1-year horizon. That would imply that you should be targeting a 99.95 percent confidence limit.

You need to define exactly what you mean by loss. At one extreme, loss could be defined as losses due to default only. At the other extreme, it could include all mark-to-market losses due to changes in credit spreads. A definition that's somewhere in the middle would be losses due to defaults and mark-to-market losses, to the extent that they're due to changes in credit migration as opposed to spreads just blowing out across the board.

I'm going to talk in a little more detail about the CreditMetrics and KMV approaches. These approaches share a common heritage, although the details are somewhat different. They had their roots in an approach that was proposed almost 30 years ago by Robert Merton. The basic idea is that defaults occur if the value of a corporation's assets falls below the value of its liabilities. The Merton approach models the asset value as a log-normal random-walk process. If at the point that the liabilities mature, the asset value is below that value, then the company is deemed to have defaulted. It's a very simple model. As it turns out, it's a little bit too simple for real-world applications, but with suitable extensions it can form the basis of a practical model. In particular, it gives a framework for capturing dependencies between multiple credits or correlations between credits.

That in a sense is what these models are all about, because if you ignore correlations and simply assume that each credit in your portfolio is independent, then what your model would tell you is that if you just add more and more credits to a portfolio, you will be able to diversify your risk all together, and you wouldn't have to hold any capital for a very large portfolio. That's clearly an unrealistic conclusion. However, it's very difficult to work directly with default correlations because defaults are rare events, and joint defaults are extremely rare events. It's very difficult to estimate joint default probabilities from historical data. The Merton structural framework provides a way of getting around this problem, because if you can link the asset distribution to the probability of default, you can use equity market data to estimate asset correlations, and that indirectly gives you a way of modeling default correlations.

Chart 4 is an attempt to illustrate the idea. I have two companies, company A and company B, and in the middle is the joint distribution of their underlying asset values. In this slide there's no correlation, so they're completely independent. This is a normal distribution. That's neither here nor there, as it turns out. If issuer A's asset value falls below a certain threshold, then issuer A is deemed to have defaulted. I can actually make this probability whatever I like simply by changing the position of that threshold. That's why I say it doesn't really matter that this happens to be a normal distribution.

Likewise, on this axis I've got issuer B's asset value, and if that falls below a certain threshold, then issuer B defaults. As you can see, the probability of both defaulting at the same time is extremely low. In Chart 5 I introduce correlation between the underlying asset values, and here I've got 50 percent correlation, which is a pretty high number, then the marginal probability of issuer A defaulting is no different than it used to be. The marginal probability of B defaulting is no different. The probability that both default has increased dramatically. This slide, therefore, illustrates that by working with asset correlations we can indirectly get at default correlations.

I put together a brief numerical illustration along these lines (Chart 6). This is looking at a portfolio of identical A-rated bonds. I've used a 25 percent asset correlation, and assumed that it's the same pair-wise correlation for each pair of issuers. We came up with that 25 percent by looking at some equity market data, and that seemed to be about in the right ballpark, and then as a point of comparison I've shown what happens if we put in zero correlation. These are assumed to be 5-year assets. I used a 5-year risk horizon because that sidesteps this whole problem of whether you should just look at defaults only as opposed to mark-to-market losses as well by making the risk horizon and the asset maturity the same date. You get rid of that issue. Because it's a 5-year horizon, I looked at capital at the 99.75th percentile as opposed to the 99.95th percentile.

This is the loss distribution for a portfolio of 100 issuers (Chart 7). In the dark blue, I've got the artificial case where there's no correlation. The light blue is the more realistic case where there's 25 percent asset correlation. You can see that once I introduce correlation on the one hand, there's a greater probability of having zero or a very low level of losses, but there's also a much greater probability of having a high or an extreme level of losses. Chart 8 shows how the capital requirements change as I add issuers to the portfolio. If there's zero correlation, then as I add more issuers, as I was saying earlier, the capital just disappears relative to the size of the portfolio. By inducing correlation, the result is that there's a systematic component to capital that doesn't shrink, and it stays there no matter how big and how well diversified you make your portfolio.

To finish off, I want to say a few words about what I see as some of the limitations of the current generation of models. One issue is that all these models to some extent rely on a variety of assumptions and parameters, and that's not to fault the models. This is simply the nature of the beast, that to come up with useful results you have to make some assumptions, and you have to estimate some parameters. However, it is very difficult when looking at credit risk on a portfolio level basis to independently verify model performance. If you're looking at a high tail-end risk measure on a credit portfolio, if you want to directly test that model, you would need hundreds or thousands of years of data, and that's just not feasible. Another issue is that the current generation of models typically treats recoveries as independent of defaults, and that's a pretty questionable assumption, particularly when you're looking at asset classes such as asset-backed securities, collateralized mortgages and so on. In fact, there is a direct link between the value of your collateral and the probability of default. A more general point is that a lot of the models that were out there were originally developed in the context of corporate bond portfolios and commercial loan portfolios, and it tends to be a bit of a stretch extending them to other asset classes that you typically find on insurers' balance sheets, such as commercial mortgages, mortgage-backed security, asset-backed collateralized mortgage obligations and so on.

A final point is the 1-year risk horizon is pretty short compared to most insurers' assets and liability profiles. There are practical problems in going longer than 1 year, but if you calibrate your model to current economic conditions, there is the potential for capital requirements that fluctuate significantly over time, which may or may not be a good thing. At ING we've taken the view that we'll use the models as they're currently constructed, but we'll parameterize the models by reference through the cycle economic data as opposed to economic data at the current point in time.

To summarize, in my view certainly these models have a lot to offer. In many respects, they’re a lot better than traditional factor-based approaches to capital. Having said that, there are some significant limitations that you need to bear in mind. Overall, my personal view is that these models are very useful in a relative sense when comparing one subportfolio against another subportfolio or when looking at changes in capital level over time. I would be inclined to take the absolute values that come out of these models with a pinch of salt. Now I'm going to hand it over to Larry.

MR. LARRY H. RUBIN: I want to go over a fairly simple model for analyzing credit risk on a portfolio of corporate credits, in this case a portfolio of credit default swaps. Since credit default swaps pay only if there's a credit event, now it can easily be applied to a portfolio of corporate credit exposures. The model itself is fairly simplistic, and I'm going to go over a model that's built using rating agency data. Then I want to go into some of the issues of using rating-agency data to analyze credit default exposure.

We begin with a pool of 100 credit default swaps, and annually for each exposure one of two things can happen. Either the credit defaults or it doesn't default. The probability of default is a function of the credit quality of the name. The model I'm going to describe here primarily uses Moody's rating to determine credit quality, although the method to determine credit quality is not dependent on Moody's. You could use S&P data. You could use KMV data. You could use internal-side credit ratings. The key to getting any good results is to try and have a method to demonstrate a good correlation between your assignment of credit quality and expected default experience. Moody's makes it easy, since they publish a lot of data on their experience.

Chart 9 shows that if a credit defaults, then the loss is equal to the notional credit less the recovery. If credit doesn't default, then one of three things can happen. Credit quality improves, credit quality remains unchanged, or credit quality deteriorates. Chart 10 is the structure of the model, and you can see it's a very simple Markov chain model. We just repeat our decision tree each year till the end of the horizon. Typically in a credit default world that would be a 5-year horizon. Now we're going to look at how we calibrate it using rating-agency data and then move onto some of the problems of rating-agency data.

If we review Moody's 30-year default rates, we see that there's never been a year where a AAA has defaulted, so we start by assuming that AAA will not default. I'm going to make one change compared with a lot of models out there. Very often, as Martin talked about factor-based models, they'll take for the credit quality a certain probability of deterioration, and then do a simulation. Are you above or below that number each and every year? Do you default or not default? When you look at published rating-agency data, you notice that there is not a smooth default rate year by year. It jumps all over the place. In many ways, the default rate itself is a stochastic variable, and that is one key thing that should be kept in mind in using rating-agency data.

If credit's rated AA, there's only been one default from 1970 to 2002. Our data could be end-point sensitive, and there could have been a default in 1969, and there could have been a default in 2003. We would get a different result than if we ignored end-point sensitivity. Therefore, the model assumes a simple distribution that 92 percent of the time the default rate for these securities would be zero, and that 8 percent of the time the default rate would be 80 basis points.

For single-A securities, there have been 3 years since 1972 in which there've been defaults. In 1982, there was a default rate of 0.26 percent; in each of 2000 and 2002, there was a default rate of 16 basis points. For this model, we assume that 10 percent of the time we had a default rate of 40 basis points; 5 percent of the time, a default rate of 50 basis points; and 85 percent of the time, the default rate was zero. For NAIC 2 securities, we find default rates best fit in a normal distribution, with the mean being the greater of the 10-year average or Moody's Expected and the standard deviation being the actual standard deviation from 1970 to 2002. Below-investment-grade securities have been shown to follow a log-normal distribution where the mean is –the greater of what Moody's says to expect, or the historical.

Now, having defined our distributions, everything's a simple Monte Carlo simulation as you've done for mortality modeling. Pick a number between zero and one. If it's less than the default rate, the portfolio defaulted. If it's greater than the default rate, the credit lived another day. If it defaulted, we look at our loss drawn from a normal distribution at 0.5, which is slightly higher than historical normal of 0.45 and a standard deviation of 0.15, which roughly replicates the historical standard deviation.

Recovery is one area where taking the simple rating-agency data again can lead you to a simplistic model. Recoveries very often differ by industry. Industrial companies, which have hard assets, seem to have much higher recovery rates than financial companies, which have soft assets, and one area where we have to make changes is actually in recovery rates by industry. If a credit doesn't default, then we determine its credit quality using Moody's 1-year transition rates, with some minor modifications to remove anomalies in the data. Generally there's a 70 percent to a 75 percent probability that the credit will maintain its ratings, and then for every three upgrades we tend to get four downgrades.

If we take this simple model and a portfolio of weighted average credits of quality A, you find that we get two defaults over 5 years, which is similar to Moody's historical average of 1.5. What we also get is significantly greater variability of results by using a stochastic loss function rather than using a static loss function.

There are also a number of problems in using rating-agency data that we want to get into. This approach is not dependent on Moody's data. It works with any data. But what are the key indicators that ratings are lagging indicators of defaults? Companies that have tended to get in trouble with credit defaults frequently load up on names that are wide for their rating class. The reason they're wide for their rating class—the classic examples being Enron and World Com—is that the market recognizes that they should have been downgraded faster than they can get to a rating-agency credit committee.

In using rating-agency data, if you have a credit in there that is trading wide to the average of credits in the same industry and the same rating, that credit is probably not the credit you think it is. You need to correct for this. Very often historical data by its nature doesn't reflect changes in the economy. The accounting fraud issues of the past 2 years were not reflected in any of Moody's historical studies, and they led to defaults that would not have been predicted, since they were not factors to be captured. The whole idea of a regime change or a change in paradigm isn't reflected.

Also, a credit that's been in its rating class for a long period of time historically has a better credit than a credit that just entered a rating class. If I have two single As, one that's been there for 2 years and one that got there last month that was downgraded, that single A that got downgraded is probably worse than the other single A. Finally, again, most working actuaries typically look at historical data and try to calibrate their models on historical data, which is what a rating-agency model does. Historical data doesn't catch regime change or changes in the economy, and it's very important to temper those looks with forward-looking data. Forward-looking data is something KMV uses, which Walid will introduce now.

MR. WALID SHINNAWI: Thank you, Martin and Larry, for providing an excellent prelude to my presentation because this will fit quite naturally. My name is Walid Shinnawi. I'm from Moody's KMV, and I'm going to cover a lot of things here. I have a lot of slides. I will try to go through some of them fairly quickly, and we'll see how much we can cover. If we don't get to cover all of them, feel free to contact Martin for the actual presentation itself, or certainly feel free to call me, and we can talk in more detail. I want to give you a brief introduction to Moody's KMV, talk a little bit about some of the challenges, as well as the opportunities, with credit portfolio management, and talk a little bit about the Moody's KMV model.

How many of you here are familiar with the KMV model? For some of you, this might be a repeat. For others, I hope you'll learn something new here. If we have some time, we'll look at a couple of case studies as practical applications, and we'll conclude with some remarks. Very quickly, for those of you who are not familiar with Moody's KMV, it was formed as a combination of Moody's and KMV, when Moody's acquired KMV back in April 2002. Historically, KMV has a fairly strong track record in building these quantitative models and tools to predict default and portfolio management, and the firm has been a pioneer in the area of credit risk management.

To give you a rough idea, we've got over 350 clients in over 40 countries spread around the world. There are many, many firms around the world, not just in North America, that are building these quantitative tools to look at credit risk management. These include looking at things from a risk management perspective, from institutional portfolio and investment management perspectives across a wide range of users. Banks were some of the early users of these tools, more recently a lot of insurance companies, and the most recent type of user would be the corporate segment, as they also have a lot of credit exposures. I wanted to give you a brief overview of who uses these tools that we're going to be talking about. Martin outlined various ways of looking at credit risk from the traditional measures to the more nontraditional measures such as the KMV measure. We're going to focus specifically on the structural type of models that have really been pioneered quite a bit by KMV, but before we get into that, we will cover some of the types of issues that one needs to think about when managing a portfolio of credit. What's the risk and return of my portfolio? Within my portfolio, what are the most attractive exposures from a risk/return perspective? Which ones are the least attractive? These are types of questions that one needs to answer as you're trying to manage your portfolios.

What's the range and likelihood of future portfolio values? Given that range and likelihood, how much equity or capital do I need to put in my portfolio? Are we earning an appropriate return on this capital? What are the major sources of concentration or diversification in my portfolio? How do I improve my portfolio performance? Which exposures should I sell or which exposures should I buy and in what quantities? I'm sure these are the types of questions many of you have thought through quite carefully, and you really need tools to allow you to quantify things so you can try to begin answering some of these questions.

The challenge we find is that credit risk is actually quite dynamic. Long gone are the days when people just bought a credit and thought, I'll just buy this bond or book this loan and sit on it and hope that I'm going to get paid at the end. What you're finding is credit risk changes a lot. It's very dynamic. It's not stable or stale as one might think by just looking at a rating. Hence, we'll see that within a particular credit rating, there are a lot of differences among the different credits and different companies. This means that our exposures need constant monitoring. We have to look constantly at what's happening within the portfolio, actively managing these portfolios, which means that now we need to actively measure not only the return but also the risk side of the equation in order to assess deals that we want to put into our portfolio.

The opportunity here is that if we are able to do this well, we can certainly enhance the performance of our portfolio, and that would lead to a more efficient use of our economic capital. That's really the ultimate goal. At the end of the day, we look at credit risk as a portfolio problem. You need to manage your portfolio of credit risk, but to manage it, you can start breaking your portfolio into its components. What is my exposure to Company XYZ? What's my exposure to Enron? What's my exposure to Rite-Aid, or you name it? Once I understand that, I need to then understand how that exposure relates to other exposures in my portfolio. We start thinking about correlations, and then, finally, concentrations. How much of each position do I have in my portfolio?

I'm going to focus on this piece first, which is the stand-alone credit risk. You can further break that into two components, recovery and default probability. What's the likelihood of my counterparty defaulting or my investment defaulting? If it does default, how much will I lose or how much will I recover? KMV has done a lot of work in this area, and it turns out that the default probability may be *the* essential piece to this whole puzzle, and we do a lot of work on that.

Before I get into default probability, I have a very quick statement on recovery or what we call loss given default. Loss given default and recovery are opposite ends of the same thing, which is an estimation of how much I could lose in the event of default. There are various ways to estimate that. Typically most people have looked at traditional look-up tables of historical recovery per, specifically for where they stand on the structure of things. Are you senior or subordinated, or where do you stand in that light? Although that's important, it turns out there is more to it than that.

As Larry correctly pointed out, there are a lot of issues around recovery, such as the industry that you're in, the economic cycle that you're in, and any relationship to the default probabilities. In fact, what you find is typically, if default probabilities rise, and your loss given default rises, that's when you get that double whammy. We built a new tool that we call LossCalc to help you better quantify the recovery. I'm not going to talk too much about that, but just wanted to let you know that there is such a tool to better estimate recovery, and basically we're using secondary market data on defaulted debt to estimate these recoveries.

Let me move on into our default probability measurement. What do I mean by the way we look at credit risk as a default probability or when I talk about a cardinal rather than an ordinal rating? We don't necessarily rate companies. Our parent company, Moody's, does that, but Moody's KMV doesn't. We look at credit risk as a probability of default. What's the likelihood of a company defaulting? It turns out that, if I look at ratings, ratings are meaningful and useful, for example, you know if AAA is better than BBB. However, you don't quite understand by how much a AAA is better than a BBB. In fact, if I look at all BBB companies, are all BBBs the same? What we find, in fact, is the market tells us that not all BBBs are alike. In fact, we can ask the next question. Are BBBs the same over time? If I have a company that is BBB today, is it the same as BBB 5 years ago? We also find that it's not, that things are changing. Credit is much more dynamic than what we see by just a rating.

What we actually like to answer is, what is the real probability of default? We're going to get into that, in a bit of detail. To do that, we found one of the key most powerful and predictive measures is incorporating market information. By market information I mean including equity market information. Why is that? We find that the equity market is very powerful and very predictive. It's very forward-looking. I'm not particularly saying that equity markets are perfect or that they're perfectly efficient, but I'm saying that often they are much more liquid than any other market, definitely more liquid than the fixed-income market. Also, there's a lot of information that's quickly reflected in stock prices, and that information reflects upon the credit quality of a company.

In addition to incorporating market information, we incorporate what we call a structural model, a structural approach. Default happens for a reason, and we need to understand what those reasons are. It is not just a pure statistical relationship, or a historical relationship. We're actually looking for a causal model as opposed to just relationships. We'll talk about this in more detail.

I'm looking at five different companies in this case. It's an old example, but it illustrates the point. Assume I'm a bond investor or I'm a loan officer looking to do some business with these companies. At first glance they appear to be very similar. They're all BBB or BBB minus, so of relatively equal credit quality, you would think. In fact, we chose them so that in this case the bond spreads are fairly similar. If I'm a bond investor, I can look at these and may choose to pick the company that gives me the best return, thinking that they're all of equal quality. However, as we delve more deeply, I would like to introduce a term here called expected default frequency, or EDF, which we'll talk about in more detail. It's the probability of default over some time horizon, in this case 1 year, and what we see is a fairly wide range of probabilities of default in these five different companies. In fact, if I focus on the two extremes, I have Cominco with a probability of default of 0.13 percent, which is similar to what we see in AA minus companies, and Avista, with a probability of default of over 7 percent, which is similar what we see in B minus. However, all these companies are BBB.

In fact, when we look at those two extremes, one is 55 times riskier than the other. It's almost like we've taken credit risk under the microscope, and we are now able to better measure credit risk by looking at these probabilities of default. The next question is, who's right? Are these right or are those right? What is the market thinking? If I take those two extreme companies and look at what happens, in this case three months later, I found that ultimately the bond market reflected the difference. The bond spread on Avista with a higher EDF widened, and you can see how ultimately the bond market reflected that Avista has more credit risk.

FROM THE FLOOR: How did you calculate the EDF?

MR. SHINNAWI: That's a perfect question and leads me into the next slide. How do we calculate this EDF? Basically we think of a firm defaulting when there's an economic reason for a firm to default, and that reason is that the firm's value as an enterprise value drops below some amount of liabilities that exerts stress on the firm. For example, if I were to go out to a bank and borrow $1 million to start a business, and a year later my business is worth $10 million, do I have a reason to default? Probably not, right? In the worst case, I can sell my business for $10 million, pay the bank the $1 million, and off I go. However, if a year later my business now is worth half a million dollars, now what happens? I still owe the bank

$1 million. I cannot pay the bank. There is no longer the ability or the incentive for me to pay it back. I just go to the bank, give them the keys to the business and walk away. This is essentially what we're saying. As long as the business value is high enough to pay the liabilities, you will not default, but if it drops below the liabilities, you will default and, hence, the likelihood of that or the probability of that is this EDF.

I'm going to delve in more detail how we actually calculate it, but before I do that, let me give you some examples. The lowest probability recalculated is 0.02 or two basis points, 0.02 percent. The highest probability is 20 percent. You can see it's a large scale so we can reflect a lot of information. Next to these numbers I've put in the rough equivalent rating, both S&P and Moody's.

Again, we're not trying to predict a rating in this case, but for most people who have never used EDFs or cannot use EDFs, if I tell you that this company has an EDF of 0.5 percent, that doesn't mean anything to you. To put things on a relative basis, we try to put these ratings next to it, and the way we've done it is to pick, let's say, S&P's BBB. If I take every single company that is BBB by S&P, calculate its EDF, and take the median of that, that's where we would place the rating. On average the median BBB would have an EDF of about 0.5 percent, just to give you a relative mix. Now I can plot the rating. I'm looking at NTL, which is the United Kingdom's number one cable television operator. Here's the rating of the company in the green line. In this case it's S&P, but this could have been anybody. It could have been Moody's. Then the red line is the EDF, the probability measure that we calculate.

For Superior Telecom. again, the EDF is rising, the agency rating is stable, downgrade, downgrade and default. Atlas Air is an air cargo company. Moody's and S&P are both upgrading, as the EDF is rising, and finally the firm defaults

There are many cases where the EDF is actually better than the rating. Conseco was another very good example. Conseco was looking fairly good. The EDF really deteriorated very fast. Rating was slow to react. Finally they had a huge downgrade but probably not enough, and then that was followed by further downgrades. France Telecom is a more recent example.

In general, you find a fairly good co-movement between bond spreads and EDFs. The bond market is fairly good at sniffing out problems, but in many cases we have found that EDFs have led bond spreads, and this is a very good example where you can see that the EDF has been rising quite some time, and then finally the bond spread blew out and followed. Why is that? Again, it's because the equity market tends to be much more liquid and much more transparent, and it's reflecting information much more readily a lot of times than the bond market.

I took all companies that are BBB, calculated their EDFs and grouped them into buckets, and in this case Bombardier is BBB. All of these are BBB, but here's

Bombardier starting out better than the best of BBBs, quickly deteriorating to among the worst 25 percent, and at this point it's worse than the worst 10 percent.

Bombardier's bond spread was stable and then blew out. By the time it blew out the first time, the EDF had shot up from about 0.1 percent to about 2 percent. That's a huge deterioration, and then it continued to blow out, now sitting at about 10 percent. This is a similar picture, but in this case what I'm showing you is EDFs, as well as bond spreads, for Bombardier as well as for the average BBB (Chart 11). If this is the typical BBB EDF, which is the dark yellow line, versus the Bombardier's EDF, you can see how at first Bombardier's EDF was looking better than the typical BBB, but it blew out around September 2001, and it continued deteriorating. You can look at the spread. The spread was looking better than the typical BBB spread, and then it blew out around September 2002. In fact, there was about a 1-year lead. The EDFs blew out in September 2001. The spread blew out in September 2002. This was a very good picture where EDFs have led even the spread in this case.

FROM THE FLOOR: Have you got any data on economic recovery, when the EDFs give false signals?

MR. SHINNAWI: Very good question. I hear this question a lot. When do EDFs give you false signals, or what if the EDF sits at 20 percent but nothing happens? The first thing to keep in mind is that, just because the company reaches an EDF of 20 percent does not necessarily mean that the company will default. What does 20 percent mean? A company has 20 percent probability of default, right? One out of five companies will default. It means four out of the five will not default. You say, well, I looked at your graph. It's 20 percent. I sold my bond and nothing happened, and I lost. Basically there are false-positive signals, right? My reaction to that is overall, if all you did was sell companies at EDFs of 20 percent and replace them with an equal return company that has a lower EDF, you should do this all day long, if you could. There is no upside in credit.

Why would I want to take the chance of having a company with an EDF of 20 percent if I can replace that with the same return with a company with an EDF of, say, 0.2 percent? What's my upside sitting at 20 percent? It's hoping that the company will not blow up. The one time that it does blow up it's going to wipe out my portfolio. In the worst-case scenario, what's going to happen? I'm going to return my principal and interest. There is no upside like in the equity markets. I would suggest to you that, yes, there might be some false positives, but in fact studies have shown that the net of these false positives still actually outperform if you follow the strategy. It's a very good question.

FROM THE FLOOR: If that's the case, then rather than selling the model, why don't you trade based on that?

MR. SHINNAWI: That was my first question when I joined KMV, because actually I come from an asset management background. I didn't want to sell the model. I wanted to trade on it. The problem is that those people who built the model and built KMV are not traders. They're not asset managers. They're academic people, researchers and quantitative people. They're very good at building models. They're not necessarily very good at trading. We have many clients who actually do this. I mean it's just a matter of doing what you're best at, right? We recognize that being a good trader or portfolio manager is beyond just building a model. You need to have market knowledge and sense, and this at the end of the day is just a model. It gives you good signals, but there are other nuances to being a good portfolio manager. We don't claim that we are the experts at that, but it's very appealing, especially for someone like me coming from an asset management background.

MR. RUBIN: Then a lot of banks know a lot of people use KMV to trade today.

MR. SHINNAWI: Right.

MR. RUBIN: Every bank is looking for that one little thing that gives them a little edge on KMV, so they can find where KMV gives a false signal that they can actually take advantage of.

MR. SHINNAWI: That's right.

FROM THE FLOOR: You have used a number of European companies in your example. When you have looked at an experience in Europe, is that different from the United States?

MR. SHINNAWI: We've done a lot of studies on our model around the world, in North America, Europe and Asia. In fact, as you saw, we have clients all over the world. We've shown that the model works around the world. It's very powerful predictor. Remember, the way the model is built is really built around an economic sense of default, and economics drive the reasons for default. At the end of the day, whether a firm actually files for bankruptcy or not, that's a different story. There are more legal and regulatory issues and so on and so forth, but from an economic perspective what drives a company to default in the United States should be very similar to what drives the ones in Europe, and should be the same as in Asia. When we've looked at the studies and the statistics, they actually confirm that, and I can certainly share with you a lot of examples as well.
Very quickly, because I'm running out of time, how do we calculate this EDF? It's basically that, if the firm's value drops below some amount of liabilities, the firm will default. The firm's value we measure as the market value of the assets. Liabilities we define as the default point. It turns out there's a third component, and that is the asset volatility or the risk of the business.

Martin alluded to some of this in his presentation. What is the firm's value? It's basically the enterprise value. We find it's very dynamic. It's very forward-looking.

It reflects deterioration and improvement well before you see that in accounting data. The problem is I can't measure that easily. I can't log onto Bloomberg or open *The Wall Street Journal* and find my asset value, but I can see the equity price. We can then measure this asset value by inferring it from the equity price using an options framework. The Black Scholes/Merton type of a framework basically says that the equity holders have an option to buy back the asset from lenders by repaying the liabilities. It's the same framework that's used to value options, except now we're valuing the entire firm.

However, instead of trying to solve for the value of the option, we know what the option is. That's the equity price. We know what our liabilities are, and that's our strike price, and we solve backwards to calculate the asset value. I'm simplifying; the actual mathematics is a lot more complicated. However, this work was pioneered and published by Bob Merton in 1974. He won his Nobel Prize based on this work. This is basically an implementation of Merton’s and other people's work. Essentially the options framework is valuable because you can now relate the asset value and the equity value in a nonlinear fashion. This is a typical hockey stick or call option payoff diagram in which, if the value of the firm exceeds the liabilities, you have upside payoff. If the value of the firm is at the liabilities, then you're basically zero. Your option is worthless. In this case now we're doing the reverse. We know what the equity value is, what the liabilities are, and we can infer the asset value. Again I'm simplifying. If you have more questions, you know, feel free to talk to me afterwards. We can talk about this in more detail, and we also have a number of White Papers on this topic.

We talk about a Merton framework, but as Martin alluded to, a Merton framework is very simplistic to be useful. KMV essentially pioneered this work by taking the skeleton of a Merton framework but enhancing it and making it a lot more sophisticated, a lot more useful. This is what we've been doing for the last 12 or so years, actually even before that.

What is the default point? It's the amount of liabilities that exerts stress on the firm. You can think of those as the short-term or current liabilities plus some amount of the long-term liabilities. I could talk about that forever, but in the interest of time I'll cover this very quickly.

To illustrate my third component, I'm looking at two different companies now. Pepsi is the company in the blue here, and Dell is in the red (Charts 12-14). We are still looking at asset value versus liabilities. If I look out here, both companies have the same asset value and the same liabilities. Which company would you say is more risky or has more credit risk?

FROM THE FLOOR: Dell.

MR. SHINNAWI: Why is that? If I don't think of volatility, and if I only think of asset value and liabilities, they look the same. I would have expected their EDFs to be the same, but when I look at the EDF, I can see Dell's EDF is much higher. Why is that? I'm also interested that, in addition to the distance between the assets and the liabilities, the fluctuations are on the asset value. We call it the asset volatility. It's a measure of the riskiness or the volatility of the company. That is our third component.

The asset volatility is a measure of the business risk of the company. We find it's directly related to the industry that the company does business in, the geography and the size of the firm.

I want to pull things together very quickly. I'm looking at a particular company with a certain asset value, this asset value will change over time, up or down, but you can estimate an expected asset value 1 year from today with some volatility around the asset value or some distribution around where that asset value is. This distribution represents the volatility of the company. I can then look at what the liabilities are, or the default point, and then measure how far I am from the asset value. What's the distance between the asset value and the default point? In this case, we're three standard deviations away. Then I can ask what's the likelihood that this asset value 1 year from today will drop to or below the default point? That is my EDF. It's the area under the curve. That's our EDF. That's how we calculate the EDF.

This is a very powerful predictor because it uses forward-looking market information. It's a structural model. There's a cause and effect. There's a reason why a company defaults. When I see an EDF increasing or deteriorating or getting better, I can quickly go through the model and understand why. What's happening in terms of these three drivers? Did the asset value go up or down? Did the liabilities go up or down? Did the volatility increase or decrease? All in an intuitive sense, what's happening to my company?

I need to take the next step to move toward a portfolio view, and to go to a portfolio view I need to have a description of the portfolio value distribution—not just the actual exposure but the entire portfolio. That's determined by how much values are going to change, and the likelihood that the changes will be happening together. It's a factor of the EDF as well as the recovery or the loss given default, but also of other things, and we'll talk about that very quickly.

Suppose now I start with something similar to what I showed you earlier. Let's call my position a BBB. It doesn't have to be a rating, we don't really care about the rating, but it makes things easy to understand. Suppose I have a particular asset with a BBB rating. What could happen over time? Certainly one option is that 1 year from now it could stay a BBB, right? However, it could also default. Here is a default/no default model. If I have a 1-year asset, at the end of 1 year my asset defaults or it doesn't. However, what if the asset has a longer maturity than 1 year? What could also happen? Besides default or no default, other things can happen.

The firm could deteriorate, not necessarily default but could get downgraded, or it could actually improve theoretically. We could theoretically go from AAA to a D, right? We can assign a probability of these events happening. Probably the most likely scenario is that it will stay BBB, but there are some fairly good probabilities of changes. Once I assign these probabilities, I can also assign a value. Suppose I start at par, and if it stays BBB, it's still at par. What if it gets upgraded? My price goes up a bit, obviously. It's not going to skyrocket. If it deteriorates, the price will go down. In the worst case, if it defaults, what do I get? I get my recovered value. In this case, if I'm assuming a recovery or loss given default of 50 percent, if it defaults, I will get 50 cents on the dollar.

Let me put this in a graphical sense. If the firm is a BBB, I get $1.00. If it's higher than or better than BBB, I get slightly above $1.00. If it defaults, I get 50 cents. If it gets downgraded, I get some value between $1.00 and 50 cents. I can then translate this into the probabilities of these events happening. If I basically take that graph and flip it, this is now what we call a value distribution. I can see the distribution of my value over time across the different probabilities and different scenarios.

Here's the trick. I can do that for one position. What about doing it for all my positions? Can I just add them up? If I have exposure A and exposure B, can I just add up their value distributions? The answer is no. When I add them up, I get a narrower distribution. You can't add them because you have to worry about correlations. That's what we were talking about earlier. You need to think of correlation. If company X deteriorates, what is the chance that company Y will also deteriorate or not deteriorate? What happens to it? I need to worry about correlations.

Obviously correlation is very important, and it's also affected by the default probability. If I have two different companies, BASF and Bayer, I can say, companies could have a bad year or a good year. There might be a low likelihood that Bayer has a very, very good year. Basically this is a low likelihood event. Bayer has enormous success, whereas BASF has also difficulties at the firm. That's one possibility.

A second possibility would be that both Bayer and BASF business values are relatively unchanged at the horizon, and so on. I can look at the various possible combinations, and if I'd assume no correlation, the value would look kind of like this slide. These are what we call concentric circles. You can think of this in a three-dimensional view like a mountain or a hill where you're standing at the top looking down. The outcome between the relationship of one company versus the other is a random event.

If there's perfect correlation, what will you see? You'd basically see a 45-degree line, and a change in one corresponds to exactly the same change in the other. We very rarely ever find a perfect correlation. A more typical correlation more elliptical where there's some amount of correlation. It's not perfect, but it's not random either. This is a fairly more typical correlation, asset correlation, between two different companies.

Remember, we're using the options pricing model to derive the asset values. We're measuring these correlations between the asset values using a factor model. We could have just measured the correlation by observing the asset values and correlating between the asset values. However, we opted not to do that but, rather, to build a factor model because we find it eliminates a lot of the noise and allows us to do correlations on companies that are not public where we cannot quickly determine the asset value. We measure the systematic risk by referring to it as the *R* squared, and then we can determine default and value correlations based on the asset correlations as well as the EDFs.

I'm going to run through this very quickly. I don't want to spend too much time on it, so as to allow some time for questions. Basically we've got 120 factors based on a number of these factors. It's a factor model that ultimately breaks things into systematic versus nonsystematic risk.

We're really interested is in this box, which is the probability of both firms defaulting. They call it the box of death. What you find is that asset correlations typically range from as low as 15 percent up to 65 percent. Like I said, we never find perfect correlation. It’s never 100 percent. We never find negative asset correlation. This is a more typical range. When it comes to default correlation, it's even smaller. You can see how it's anywhere from 2 percent to 15 percent. It's a much smaller event.

As we said, that credit quality is very dynamic. Correlation is very dynamic as well. Why do we care about correlation? Is it important? Let's illustrate this through an example. Suppose I'm looking at two different companies. I'm looking at Compaq and International Specialty Products. Let's suppose I'm looking at exactly the same structure, a 1-year, $50 million deal, senior, unsecured, fully drawn facility. I'm looking at these two companies. They're paying me the same amount of fees. If I look at the EDFs, it looks like Compaq has the better EDF. Which company would I choose at this point?

FROM THE FLOOR: Compaq.

MR. SHINNAWI: I’d choose Compaq if I'm looking at things on a stand-alone basis, but if I'm adding things to my portfolio, do I have enough information yet? I need to know more information. What if I look at the unexpected loss or one standard deviation around the expected loss? Compaq still looks better. Do I have enough information? Probably not. What I really need to know is the contribution to my portfolio, to my portfolio risk. On a stand-alone basis Compaq looks better, but within my portfolio things may be different. What if now we look at this and we see that the *R* squared for Compaq is 57 percent? Now things may change. Now my risk contribution for Compaq is actually higher. It's possible that within my portfolio I have already too much Compaq or maybe I have too many other companies in the high-tech sector, so that adding one more high-tech company is adding risk to my portfolio, despite the fact that Compaq on a stand-alone basis looks better. I can see now my return on capital on Compaq is a lot lower. This is why we need to worry about correlation, to put things in perspective.

When we talk about the role of capital, think of capital as a reserve to absorb losses. It can be allocated to a portfolio based on some desired safety of liabilities issued against that portfolio, and the objective is to have sufficient capital so that the risk on these liabilities is equivalent to some kind of a debt rating that you're targeting. Let's say I want to use an AA rating for my firm, and suppose an AA corresponds to about five basis points of default risk. As a guide then, I need to have enough capital to guarantee solvency 99.95 percent of the time. If I use this as an example, and this is my loss distribution, to achieve an AA rating, which is right here, I can now read through and see how much capital I need. That should cover me up to five basis points of loss. That's how I could use these tools to come up with capital calculation.

Chart 15 is a view of your portfolio. Every single dot here represents a specific position within your portfolio. If you click on the dot in the model that we have, it tells you the exact name and amount and quantity and so on. This red x represents the position of your entire portfolio. We're looking at things here from a risk versus return perspective. Certainly the portfolio looks better than each individual position, because of diversification of the correlations. This green line represents your efficient frontier, and then the various dots on the line represent various positions on that frontier, and you can kind of run optimizations to see how you can achieve the efficient frontier. The idea, of course, is that you can enhance things by either improving your return or reducing your risk, and if you do this, you're shifting up to the left.

We talk about mispricing. We can figure out how much spread I need for a particular exposure. In a hypothetical example, Ford stands relative to where the required spread should be. I need more spread or maybe I need to lighten up on Ford compared with the next example, where I'm looking at Carlsberg. I can actually earn less spread and still be okay, or it tells me how much I can load up in my portfolio at that given spread and still be okay. I can do those types of analyses. The idea again is that performance is a gradual process, and you do this over and over again. You should think of the mispricing graph that I showed you as your road map to performance of your portfolio. Because of the time I'm going to skip through the case studies, but the case studies pretty much illustrate the same points that I've talked about in terms of correlations and diversifications.

MR. LE ROUX: Are there any questions?

MR. JOHN MARCSIK: I'm from AEGON. You mentioned that you use option pricing to back out the value of the assets, but it seems that you left out one variable, because you didn't mention the time horizon that you use there.

MR. SHINNAWI: We assume a perpetuity. We assume that the auction is good forever, basically. That's the underlying assumption. We're assuming basically an American option, I believe, where you can exercise it at any time in the future.

FROM THE FLOOR: Yes, Mr. Shinnawi, your model gives you these probabilities of default. Have you then tested whether these companies with a given default probability default with that probability?

MR. SHINNAWI: That's a very good question. Unfortunately I don't have the pictures here to show you. I have them on my computer, and if you'd like, I can send them to you. We've done a lot of studies to test these models. In fact there's a paper, if you're interested, called "Modeling Default Risk," which is available on our Web site. If you go to kmv.com, there's a research section. There are several White Papers, in particular "Modeling Default Risk," which goes into the methodology, but the end of it has a lot of these performance studies. We've tested a number of things. First, we're calculating these probabilities of default. Are these correct? We look at actual defaults and see what happens. We found that the actual defaults were very, very close to what we have predicted, and we've done this over a number of time horizons.

The other thing we've done is these par curves to say, all I'm using is one system to look at good names versus bad names. How would I have done? For example, if I have a portfolio of a hundred names, rank them from best to worst by rating agency, eliminate defaults by eliminating the worst rated ones, and see how you would have done versus ranking them again by EDFs, we found you will always outperform using EDFs. A particular test that was published in *Risk* magazine did very similar things. It took various names within a particular rating bucket, and then tracked default over 6 months, 1 year, up to 5 years, and saw the performance. Again, it performed quite well. Unfortunately, I don't have the slides to show you, but I can send them to you if you'd like. Give me your e-mail, and I'll forward those to you.

Chart 1

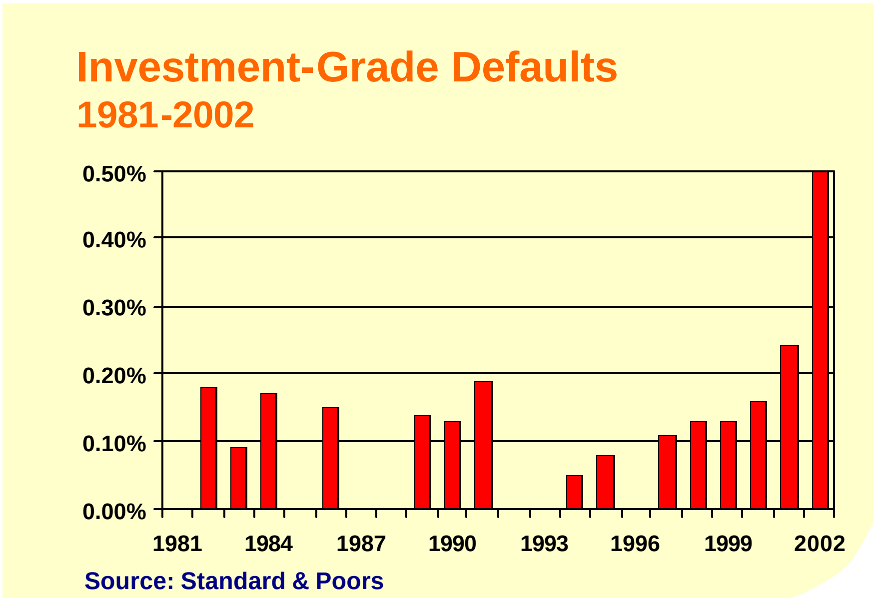

4

Chart 2

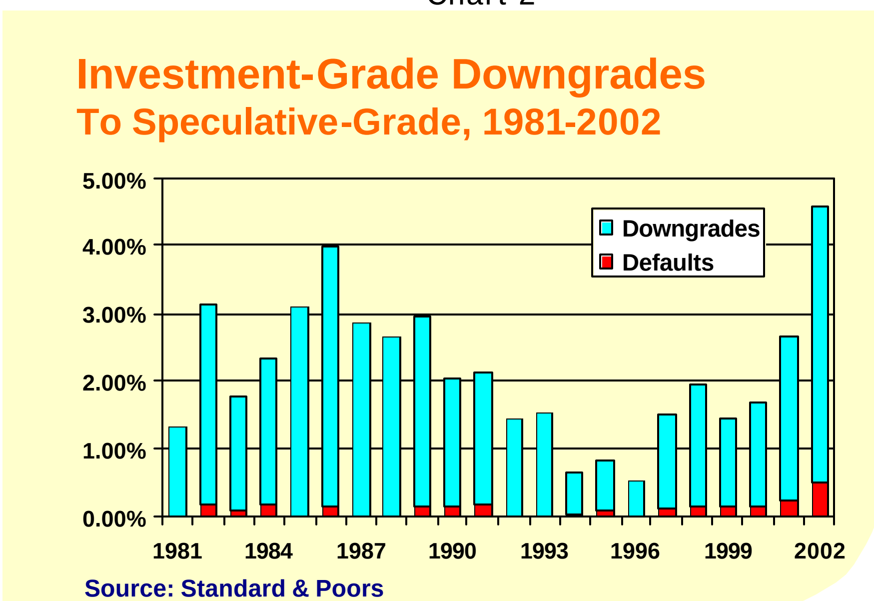

5

Chart 3

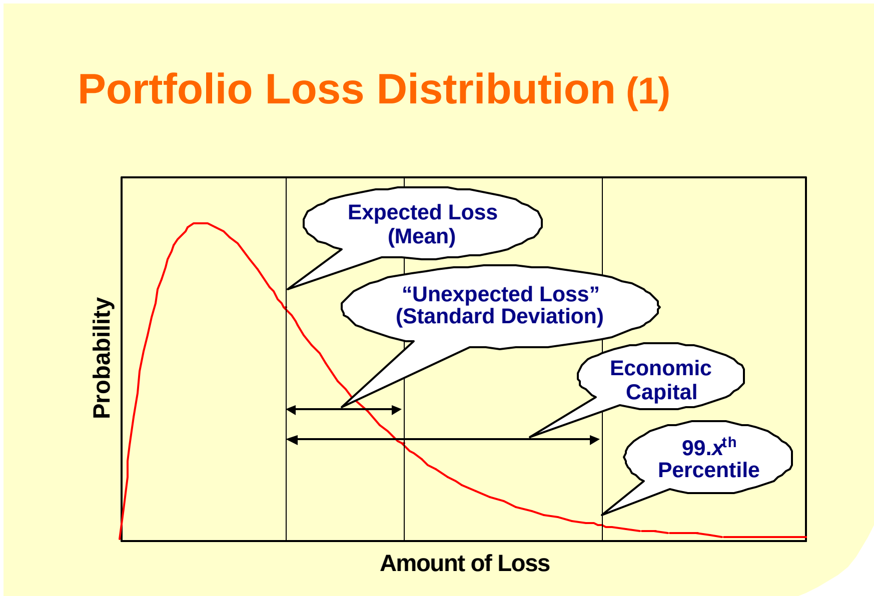

9

Chart 4

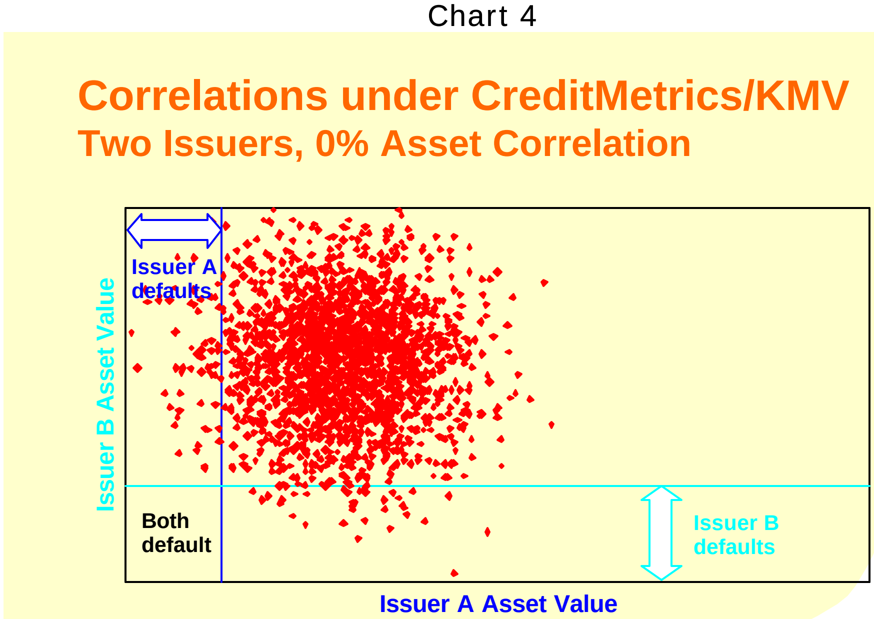

14

Chart 5

Correlations under CreditMetrics/KMV

Two Issuers, 50% Asset Correlation

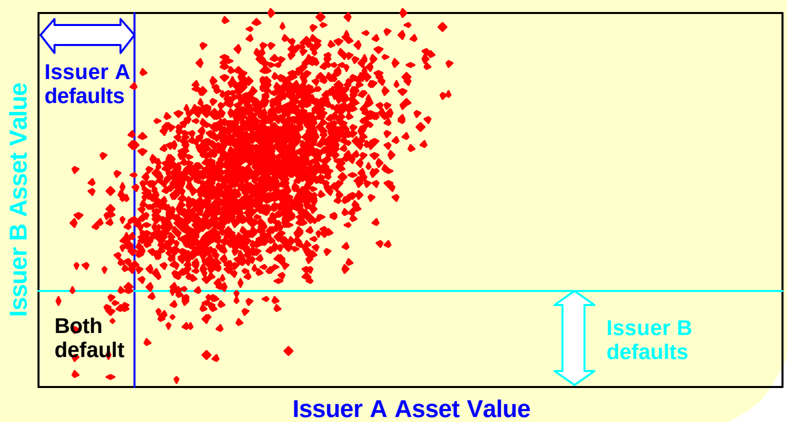

15

Chart 6

Numerical Example

Required Capital: Corporate Bond Portfolio

- **Portfolio of 5-year A-rated corporate bonds**
  - **Identical notional amounts**
- **5-year default probability 0.55% for each issuer**
  - **Based on S&P data, 1981-2001**
- **45% recovery on default**
  - **Historical average for senior unsecured corporate bonds**
- **25% asset correlation for each pair of issuers**
  - **Also looked at 0% correlation for comparison**
- **Capital set at 99.75th percentile, 5-year horizon**
  - **No need to simulate mark-to-market gains or losses at intermediate horizons**

16

Chart 7

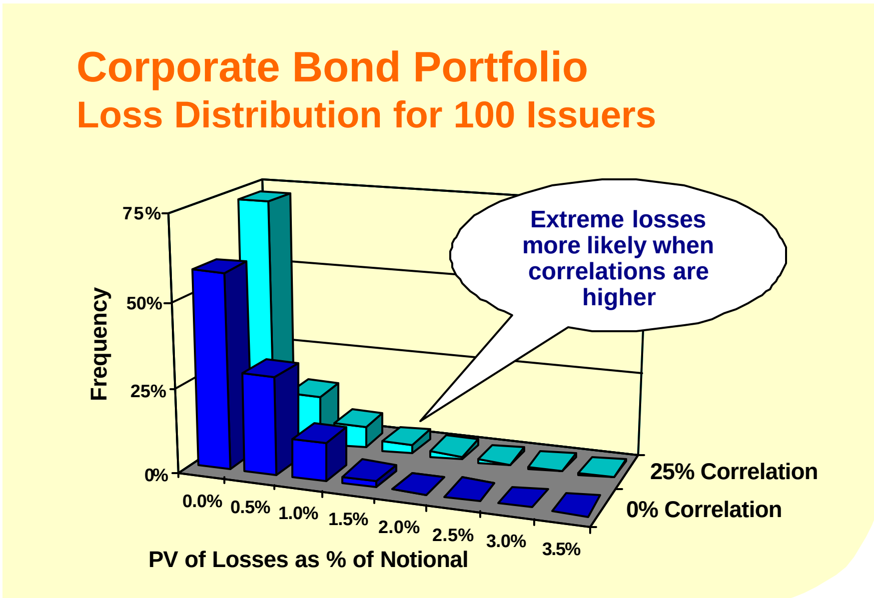

17

Chart 8

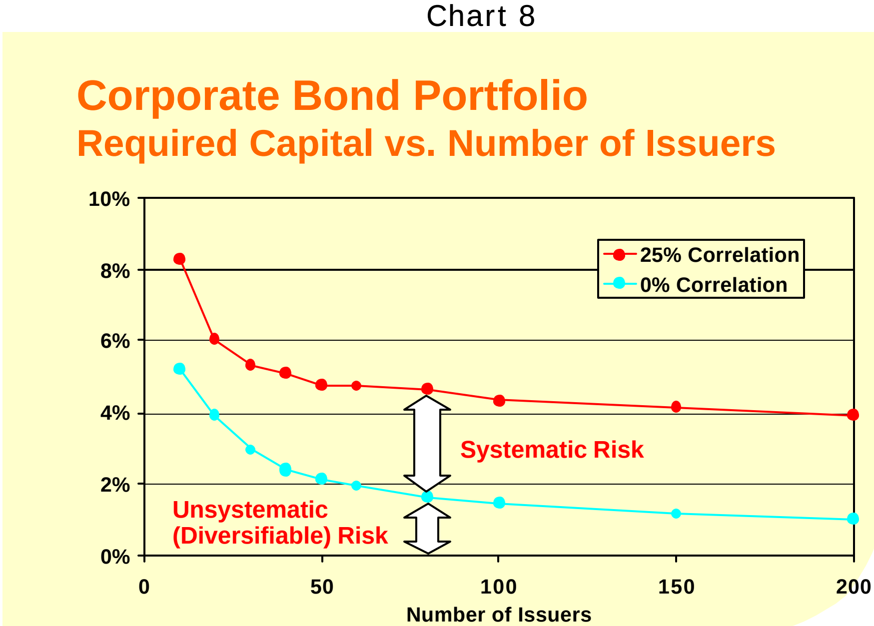

18

Chart 9

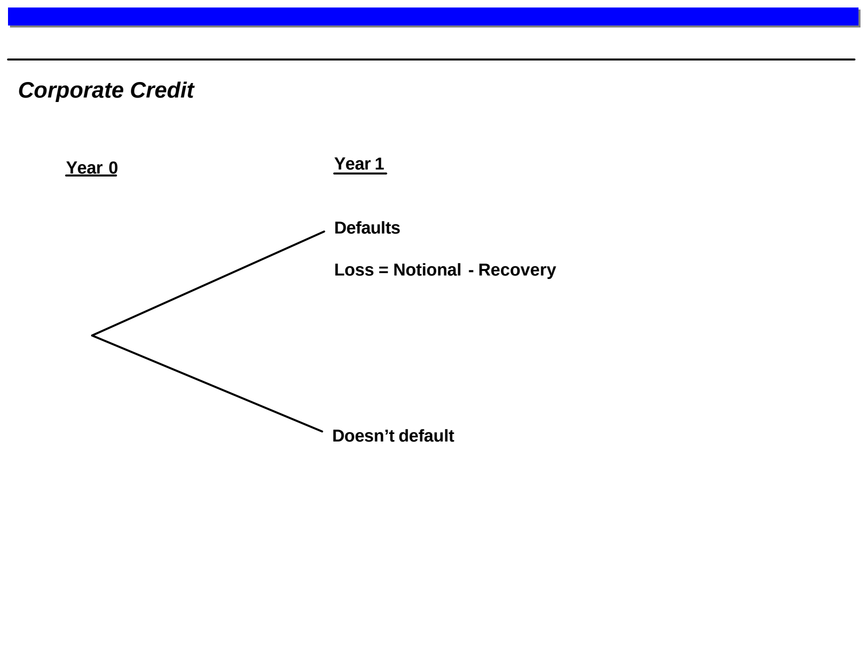

2

Chart 10

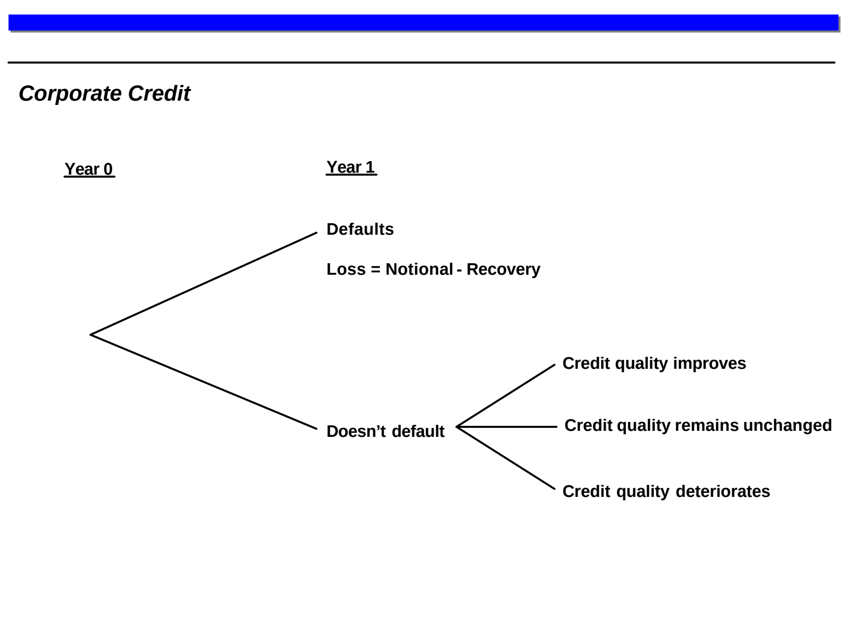

3

Chart 11

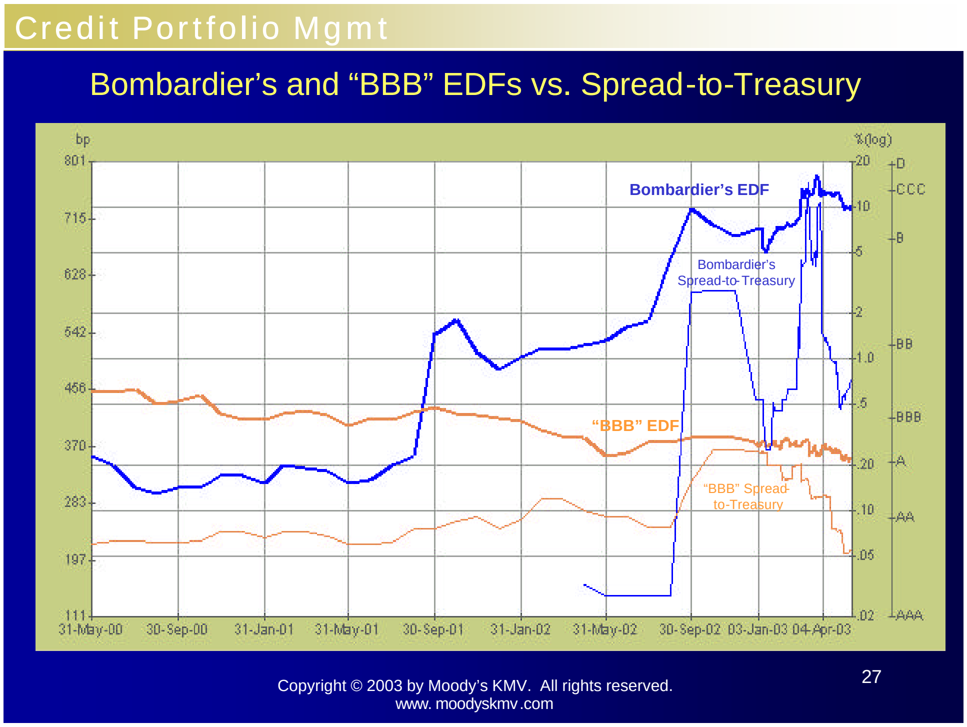

Chart 12

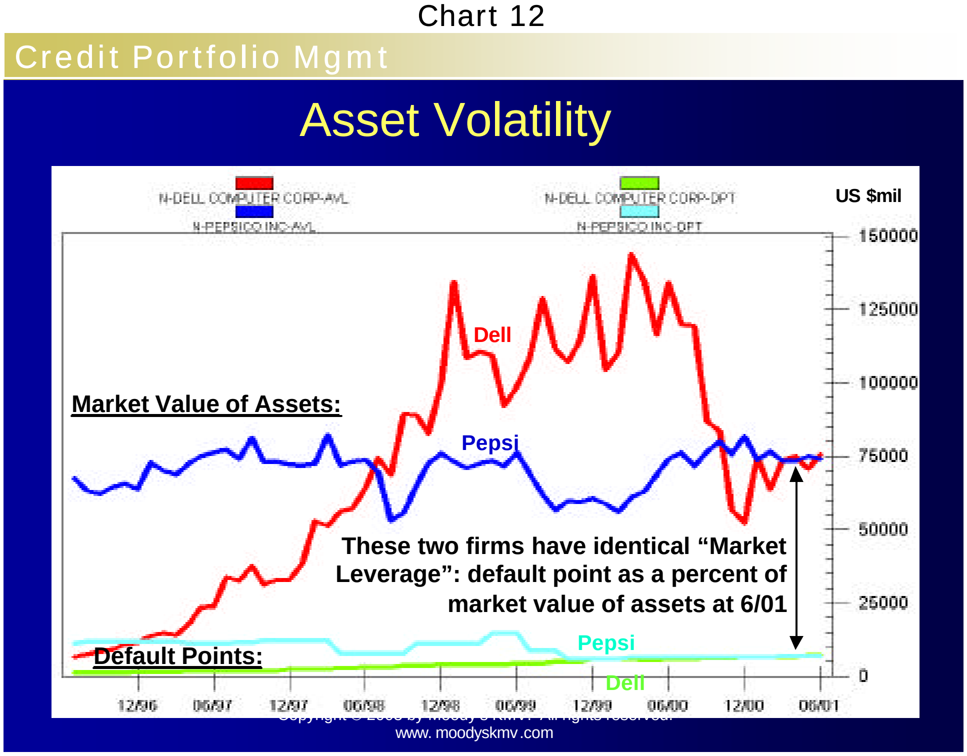

Chart 13

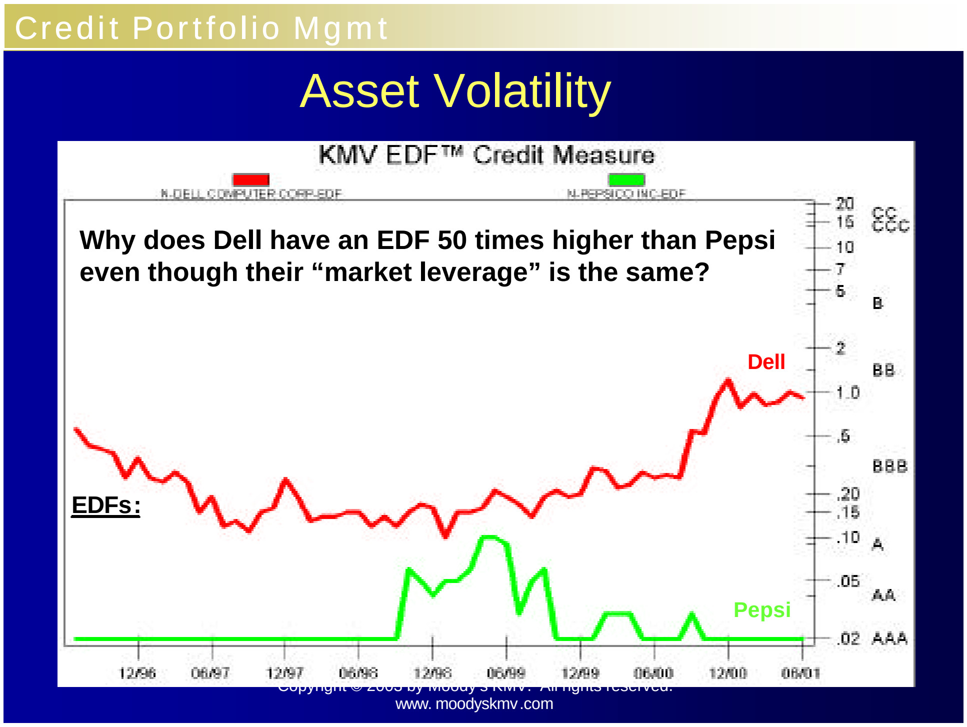

Chart 14

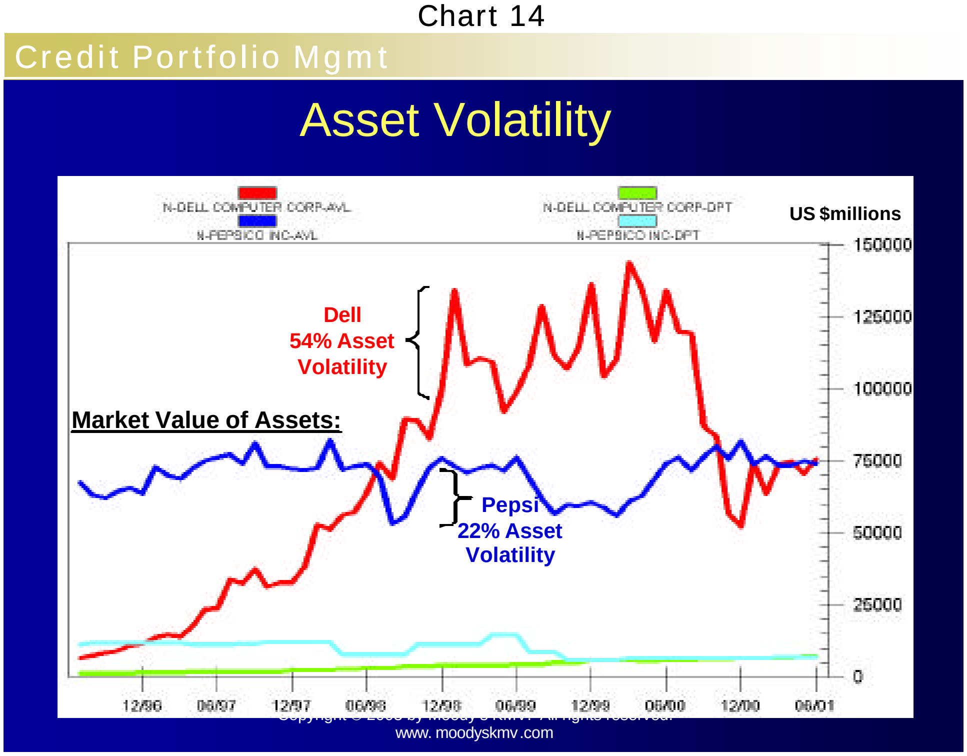

Chart 15

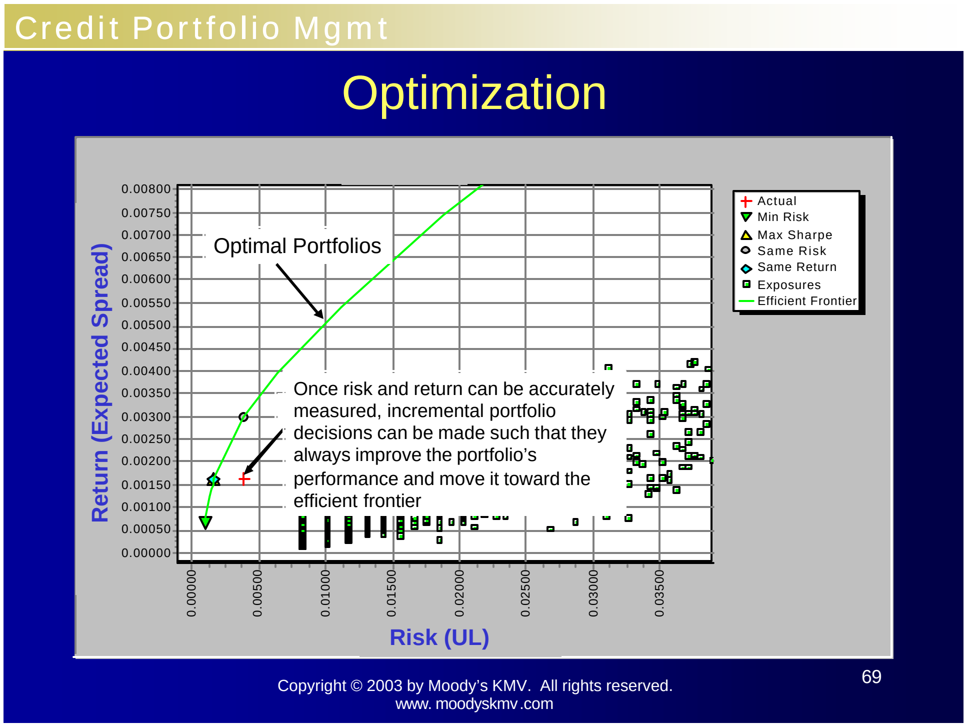
Credit Portfolio Mgmt
Optimization
Return (Expected Spread)
0.00800
0.00750
0.00700
0.00650
0.00600
0.00550
0.00500
0.00450
0.00400
0.00350
0.00300
0.00250
0.00200
0.00150
0.00100
0.00050
0.00000
0.00000
0.00500
0.01000
0.01500
0.02000
0.02500
0.03000
0.03500
Risk (UL)
Optimal Portfolios
Once risk and return can be accurately measured, incremental portfolio decisions can be made such that they always improve the portfolio's performance and move it toward the efficient frontier
Actual
Min Risk
Max Sharpe
Same Risk
Same Return
Exposures
Efficient Frontier
Copyright © 2003 by Moody's KMV. All rights reserved.
www. moodyskmv.com
69