RECORD, Volume 30, No. 3*

Annual Meeting and Exhibit
New York, NY
October 24-27, 2004

Session 14 PD
Integrating and Aggregating Risks

Track: Risk Management

Moderator: Francis P. Sabatini

Panelist: Ugur Koyluoglu

Summary: Risk integration and aggregation are key issues right now for many companies as there is greater demand for transparency regarding risk exposures. There is a need to provide a company-wide view of risk across all elements: interest rate, market, policyholder behavior, credit and operational. Implementing risk integration and aggregation provides a powerful view of enterprise risk and the benefits of diversification on the total company risk exposure. The panelists discuss approaches for measuring risk across risk elements on a consistent basis and discuss the methods for aggregating results and measuring diversification.

MR. FRANCIS P. SABATINI: We'll start with Ugur Koyluoglu. Koyluoglu is a Ph.D. from Princeton in operations research. He spent four to five years teaching applied math and engineering. He has spent eight years now in consulting with Mercer Oliver Wyman, both on the banking and insurance sides. I'm with Ernst & Young, and I'll follow.

MR. UGUR KOYLUOGLU: I would like to present to you some challenges of integrating and aggregating risks and the applications for which you can use the aggregated view. I will have examples from life insurance and banking. I spent more of my career in banking, and there are some interesting lessons that I would like to share with you.

If you think about the uses of the output of probabilistic risk aggregation, from an actuarial point of view, the starting point is the calculation of diversification

Note: All handout materials are available through the link on the table of contents.

benefits. There are many applications with corporate finance. You can think about capital attribution. Then you can reflect that to risk-adjusted performance measurement. You can think of identifying key risks and introducing risk control and limit setting around them. Obviously you can take the aggregated view to your risk-based pricing evaluation and in thinking about how to use the integrated view with the accounting financials in terms of understanding earnings volatility and its key drivers. Last, but not the least, you can check whether you have an optimum hedging program for interest rate risk-taking if you're thinking about everything from an aggregated point of view versus a single product-specific optimization.

Our time is limited, and because I would like to reach most of you, I'll try to be nontechnical. Nevertheless, I'll go deeper in terms of explaining the concepts for architectural designs for risk aggregation.

What I mean by that is that you can either start with risk drivers—these could be market economic variables or a set of key risk drivers at your institution—and try to aggregate all the risks up, or you can start from a component approach in which the experience in each risk type would try to identify the total risk distribution for each, and now you can think about how to bring all of them together. So, those are the two architectural designs that I'll cover today. As I explain these to you, I would like to focus on old-school practical and numerical challenges.

So, what do we mean by diversification benefits and at what levels could you really calculate the benefits? You can think about the diversification benefits within a risk type and within a product. For example, you could consider the diversification benefits on the investment side in the corporate bond portfolio. Or you could consider across a risk type but still within a product group. For example, we can focus our attention on fixed annuities and then think about all the risks in fixed annuities and aggregate them together—the interest rate risk, the credit risk and mortality risk. Another way of thinking about diversification benefits is within a risk type but across product groups. Again, an application is the optimization of the hedging program. You could think about diversification of interest rate risk across fixed annuities, institutional products and other interest-sensitive life. Last but not least, you could consider across everything—risk types and product groups—and the following approaches could be applied at all levels. You might prefer a risk-driver approach in aggregation at some level and a component approach at another level. This much depends on the variability and how you manage these risks.

There are several sessions on risk management today, and I'll assume that somehow you are happy with your risk identification and characterization of risk distribution. The problem that we will focus on is aggregating those. In Koyluoglu Slide 4, I divided the risks into pieces for simplification. You could follow any risk taxonomy. For credit risk, the investment side works with its own modeling. They have mark-to-market (MTM) portfolio management or risk metrics, whatever, and it gives you a probability distribution for the current risk-taking. For asset risks, again, there is some modeling going on. For mortality risk and asset/liability management (ALM), it could be in an integrated way. You might have your own system coming up with the distributions of each. For the other risks, there might be some consideration or no consideration in your organization, but I have seen lately that there's an increase in giving attention to operational risk modeling, so you might have a model for that. Business risk modeling can also be integrated to the ALM risk modeling through lapse rates and consideration of volatility of margins in the future.

Say you have done all your analysis, and you are happy with that. The next problem is aggregation of those to come up with a company view, a product view or just a single risk-type view. Herein, we need more information, and that is to describe the dependencies or correlations. I use the term "dependencies" just to generalize correlations because, if you think about these risk types, the dependencies could be nonlinear, and correlations here and within my speech just focus on linear correlations. So you bring all of them together, and you could come up with the aggregate view. From the aggregate view, again, you can focus your attention to capital, any percentile such as earnings volatility.

If you think about learning from other institutions, several banks have already done this, and they are disclosing their results to the investors and the public. The chart in Koyluoglu Slide 5 is from J.P. Morgan Chase's annual report. They have done the analysis, and they post their stand-alone capital requirements for each risk type, as well as the diversification benefits. Apparently they think that they need $13.1 billion of economic capital for credit risk, $4.5 billion for market risk, $3.5 billion for operational, $1.7 billion for business and $5.4 billion for their private equity investments.

Adding these up, without considering correlations or dependencies or assuming that all of these are 100 percent correlated, they ended up with an economic risk capital of $28.2 billion. Then they bring in their expertise for risk aggregation, interaction of dependencies, correlations and diversification benefits. As a result of that, they have calculated their diversification effect to be $5.1 billion. So, the total capital needs are then $23.1 billion. There are other things in the calculations that I'll skip because I want to be consistent with what Citigroup reports.

On the right side of the slide, you see a presentation by Citigroup's chief financial officer (CFO). They more or less defined the risks in a similar way, and they have calculated their capital needs. Credit risk is the largest. Then comes the market, then operational, then insurance—they have a small bit of it—and they realize that their diversification benefit will be about 10 percent of the stand-alone sum. So, if you add the stand-alone capital that they believe they need, that adds up to $51.9 billion, and they believe the diversification effect is $5.2 billion.

This is a very simple view. In the end, investors and rating agencies, in addition to regulators, would like to see this view. However, there's a lot of effort behind this. We certainly underestimate how much time and effort is needed to get these results and be comfortable enough to disclose them. I believe that in both institutions, at least 100 man-years—probably around 300 man-years—of effort are needed to reach this stage. In that effort, obviously, calculations are important and the results are important. But the discussions are priceless, because you embed this into the culture of the organization and then reflect these measures to performance measurement.

If you think about the risk aggregation behind the scenes, there is a lot of risk aggregation going on. Let's look at credit risk. They have multiple transactions with just one customer and then also some guarantees to the group. Then there are several products, credit lines, leases and all sorts of other commitments. They also have exposure in various parts of the organization throughout the world. So, a lot of risk aggregation is going on within one risk type.

Then across risk types, they already introduced the correlation structure. We will talk in a bit about how you might want to do that in life insurance. There are serious challenges in parameter estimation. They also do not stop after they disclose the economic capital to the public. Before that, they have already embedded this into their understanding of risk management. So, in the example, I wanted to illustrate the questions that they are looking to answer and how they are using such a system. How much capital is needed? Do we have excess capital? Where is capital attributed in your organization, and what is the return on that capital? What do shareholders require for risk-taking, and what is the value of a business?

In the example, business units A, B, C and D are capital consumers, and A and B are apparently heavy capital consumers. It could be because of large volume. It could be because of large volatility or high correlation in the rest of the portfolio. When we look at the risk-adjusted returns, apparently business units A and D are very profitable in a risk-adjusted sense, because they are yielding a risk-adjusted return of 35 percent and 40 percent versus a hurdle rate of 15. So, these businesses are really creating value. In the example, we just calculated value using a dividend discount model. What business unit A does is that when you put in capital of $1.3 billion, it creates value, and it's really worth $3.9 billion.

Leaving what banks do aside, let's go back to our topic, which is the challenges and design issues in terms of risk aggregation. The first challenge is consistency of risk measurement before aggregation. Are we measuring all the risks that we are pulling together using the same units? If you are using value on one side and loss on the other side—economic value or loss on one side or accounting value or loss on the other side—it is mixing things up. How can we achieve consistency as much as possible?

Time horizon is important. Are we looking at a very short time horizon? Are we looking at one-year or multi-year horizons? If so, how will you bring all these views together? Then there is the common risk metric. On the banking side, confidence intervals seem to become the common risk metric. On the life side, as far as I can follow from the recent publications, expected tail loss, or conditional tail estimate (CTE), has been receiving a lot of attention and attraction. We need to think about the consistency issues before aggregation.

Then once you start thinking about aggregation, the biggest challenge that you need to meet is, how would you like to model dependent structure? If it is correlation matrix, how will we estimate the correlations in the matrix? If it is factor-based modeling, will we be happy with all the regressions that we have that explain the loss or value from the macroeconomic factors that are driving it? If you don't want to use linear measures, though, can we use copulas or things like that? Mathematically, yes, but which copula? These are challenges, and I'll emphasize that one the most. Then say you decide to follow a simple approach, correlation metrics, when your correlations are good enough. Let's do a set of runs, and then we'll see how to enhance the modeling.

The next challenge is picking out the right historical data that you feel is relevant to calculate these correlations. These correlations are highly volatile all the time. You might need to introduce your expert judgment and perhaps have a lot of discussions internally to agree on these dependencies—use Delphi method or introduce parameter error, and then there are all these non-stationary dependencies.

As we approach this, I would highly recommend that you think about a starting point and calculate what you are after and then test the sensitivity of the results or decisions or other discussions with respect to the assumptions. Then you identify what is important, and you can revisit your calculations accordingly.

As for the aggregation technique, in the introduction I mentioned risk component versus risk driver. In terms of getting it done, you are the masters of using all the probabilistic techniques or Monte Carlo simulation or numeric integration to make it work, so I won't go into the details of how to finally put it in place. Let's say that you put it in place, you came out with the aggregated view, and you did it like J.P. Morgan Chase or Citigroup. You calculated diversification methods—$5.2 billion for Citigroup and $5.1 billion for J.P. Morgan Chase. How will you allocate this back to risk types or back to products? Do we want to allocate it back to everybody, or do we want to keep it at the corporate level? What are the corporate finance implications of doing so? How can that be reflected to risk-based performance? If you allocate, will we use something like a 10 percent haircut to everybody? Will we think about the marginal contributions? Will we introduce Shapely value? I don't know how familiar you are with these topics, but I wanted to put them forward as challenges.

The next slide (Koyluoglu Slide 7) is the summary of risk component approach versus risk driver approach. I used to be an engineer, so I prefer just two very simple designs. In the risk component approach, basically you would put in the "best of breed" methodology that you could for each risk type, and then you would introduce the dependencies on top of that. In terms of the structure, your columns will be pretty good, but the beam that connects all of them, or the roof, wouldn't be that good because it's very difficult to introduce a dependence structure (correlations) among the component risks. In the risk driver approach, you will identify a set of drivers. These are generic drivers, such as interest rates, equity returns and mortality. You can also go look at macroeconomic indicators, such as gross domestic product (GDP), unemployment rate, consumer price index (CPI) or whatever. Then, through regressions, you will bring the dependencies together. Here you will have a much stronger beam or roof, but perhaps your drivers are not really a universal set, so the columns might not be as strong as you think.

Let's look at an example (Koyluoglu Slide 8). Frank has some more interesting examples specifically dedicated to life insurance. This is for an annuity business. This was from a client, and they had models in place. They had a different credit model for the bond portfolio, another one for real estate and then another one for mortgage-backeds. They got all the output, and they wanted to bring all of these together. Apparently, because all of these are managed by different people in different parts of the organization, they prefer the risk component approach. The challenge was to introduce the correlation matrix across these four to come up with the aggregate view for the earnings. Obviously, leaving all the accounting part and smoothing part aside, this will give you indications about the drivers of earnings volatility. It won't match exactly the earnings predictions that you might have from accounting systems.

The main challenge was to introduce the correlation matrix. It's not that easy. If you think about it, these could be the correlations between interest rates, equity, mortality, credit defaults of different products, different ratings. For some of these, you would have historical data, but we can't estimate correlations directly from historical data. For others, you might need to just introduce your intuition. Obviously you could bring in your internal data. You could bring in your expert opinion, but you shouldn't expect something that you'll feel very comfortable with in the first iteration.

I've started to call correlation, or dependence modeling, "define your comfort zone." You should start from somewhere and calculate your output that will impact your decisions. Then start testing the sensitivity of the output with expected assumptions. On Koyluoglu Slide 10 you see sensitivity of the results. This client defined economic capital at the 95th percentile, and the sensitivity of that with respect to the correlations and the metrics. We change the correlations one by one and test the sensitivity of the results to identify the key correlations so that we would focus our attention to those. As you can see, there are a few that drive the results. You can then focus your attention once you identify those.

In terms of the risk driver approach (Koyluoglu Slide 11), first you need to identify all risk drivers. These would be interest rates (if you are a global player, exchange rates), the equity market indices, GDP, default rate, etcetera. You need to characterize the distributions and the correlation dependencies across them. Then you need to introduce the sensitivities of credit risk, ALM, equity and other risks with respect to these drivers. Finally, you can aggregate them together. But it all starts with the modeling of risk factors. These risk factors in the example are lapse behavior, yield curves, corporate bonds, credit defaults and equity returns. Then you are building the projections of balance sheet and profit and loss (P&L) just out of those.

Again, once risk drivers are identified, you also need to do some modeling. You can think of starting with historical data. You can introduce all the implied volatilities. Depending on your time horizon, you might need to reflect your expert opinion in the modeling of risk drivers as well. Here (Koyluoglu Slide 13), we simplified a couple of risk drivers: 10-year Treasury rate, the spec grade default rate, Standard & Poor's (S&P) 500 Index and monthly index return. Its modeling is illustrated from historical data, but we should also look at the implied volatilities. Once you have modeled all of those risk drivers and the dependencies or correlations across them, you can simulate a set of consistent scenarios, and under these scenarios you can run your individual models for aggregation. We follow this technique as well, just to cross-check the component approach. The same set of scenarios is going to the portfolio model on the side, and it was producing results so we can really aggregate them.

These are different ways of approaching the problem, but these are not the only two ways of doing it. Another way of designing the aggregation is simplifying everything to a few factors. That's the way banking regulation handles risk aggregation. I started to see similar approaches discussed in future insurance regulation, especially in FSA discussions. Here we are just looking at the credit risk side.

So, basically, in terms of aggregation, you define a single factor, and that defines your correlations across risks within that risk type. So, a single factor could be the world economy. Let's say that there is such a factor. Then the world economy, the systemic factor, is not doing well. Old bonds will be distressed. As the systemic factor is doing well, you are in favorable economic conditions. If it is doing poorly, you are in unfavorable economic conditions. Link the output in which you are interested—that is, the conditional default rate in your bond portfolio—to that systemic factor. So, whenever a systemic factor is bad and you can define the confidence level, then you can figure out what will happen to your default rate.

The whole banking regulation Basel 2 is resting on this simplified approach and the description of the linkage here. Linkage of conditional default rate to the systemic factor is through what they call a Merton Model, and it assumes that a firm will default when value of assets is less than the value of liabilities. It's not that complicated. You just define one risk factor, and that risk factor drives the whole risk. The aggregation is done automatically because you just stress the risk factor, and then you look at how different ratings are responding to that. Automatically, you have aggregated everything up.

This finalizes the design issues that I wanted to introduce. I want to emphasize three points. The first one is dependence measure. Do you want to use something simple or more complicated? As it becomes more complicated, you definitely need to do sensitivity testing. The second one is the consistency problem. You have short-term and long-term risks. How will you bring all of them together? As you think about long term and short term, there are implications of the correlation estimation because correlations for a one-year time horizon are different from those for a 20-year time horizon. The last one is that you can think about different approaches. I believe these are secondary compared to the other two. As you study the sensitivity of the results, you can really see those all by yourself as well. Frank would like to go through some practical examples.

MR. FRANCIS P. SABATINI: I think you'll find that there are a lot of parallels in what we both have to say today, with slightly different perspectives, so I think it will be particularly useful. I want to talk a little about why this is an important topic. The world we live in is constantly changing, and we've lived through a pretty difficult period over the past several years. We've had low interest rates. I'm sure that each of you, as part of your organizations, has had to think about the implications for the organization if interest rates continue to stay low. That's a continuing issue.

We've had a few years of a dramatic bear market and a little bit of a recovery, and we're going sideways these days. With the proliferation of equity market-related products that we sell, volatility of income results and the exposure that we have to guarantees, this becomes an ongoing problem. Not too long ago—I think three years ago—we experienced one of those credit events that we seem to experience every several years. We can't forget that they'll probably happen again. We've even extended the guarantees that we offer to the marketplace. All of that has created a greater desire for transparency by external constituents in terms of the balance sheet and the amount of risk that the organizations are taking.

I was intrigued by the disclosures that Morgan Chase and Citigroup are making because I think that that's where we're headed. The investors, the rating agencies and the analysts want to know how much risk these organizations are taking. They know that the exposure from company to company is different, or at least they think it is, and they behave that way. If you listen in on a few analysts' calls and the amount of Q&A around risk issues these days, you'll see perception driving analysts' behavior. It may have nothing to do with reality, but it certainly has a lot to do with their perception of risk that a company may have. So this desire for greater balance sheet transparency as it relates to risk, I think, will drive change. It will drive the actuarial profession toward developing techniques for integrating and aggregating risk and helping the external world, as well as management, to understand how much risk is really on your balance sheet.

One of the interesting things to take away from this Citigroup and Morgan Chase illustration is the relative size of the different risk components. It at least allows those organizations to decide whether or not they have too much credit risk relative to their overall risk exposure. The other interesting thing—you'll see it throughout my talk—is the idea of the correlation benefit. I actually believe—it's more intuition—that insurance companies have more risk diversification on their balance sheets than other financial institutions.

So finally, I want to get to the last point, and that is management's desire to understand how much risk they really have. I've been involved in many situations with management. I was with a CEO. We had just gone through a presentation around risk exposure on a particular product line. He walked out of the meeting and said, "That was a great presentation, but I really don't know how to deal with it because we have all these other businesses. I don't know how this exposure that we've just talked about relates to the total exposure in the company." I've been with a CFO recently who said that they just implemented a hedge program around some guarantees that they had been issuing. He said to me, "We really didn't have to implement that program, because the exposure relative to the size of our balance sheet is relatively small. We could absorb the risk exposure. But we had to implement the program because the external audiences were demanding it of us, and we couldn't defend our position in terms of balance sheet diversification." I'd be curious to hear the extent to which your organizations are starting to think that way. This is the driving need, and this is why it will become more important.

Redefining the game means getting the external constituents refocused on a broader balance sheet rather than on specific issues. For example, a rating agency independently deciding that a particular product line has too much risk limits your exposure to that product or adds capital charges for that product without any recognition of the diversification that you might have across the balance sheet. It creates pressure on management to respond, and it leads to decisions that are made that may not be necessary. It invites a discussion with the external world that's out of context. The solution clearly—I'm sort of overstating my case at this point—is to put the issue in perspective and create a balanced discussion of risk within the organization and with the external world.

What do you need to know? There are questions that need to be answered. How much risk is on the balance sheet, not only across product but also across risk element? How are the major risk elements represented? Are the measurements equivalent? I think we will have succeeded on the day we can say that we can actually compare a dollar of credit risk to a dollar of C-3 risk or ALM risk exposure, because those are the two risk elements that drive most of our capital need, other than maybe operational risk, which is a much different discussion.

How much benefit do I get from balance sheet diversification? Do I have too much or too little exposure? I'll touch on it a little in this speech, but the whole idea of risk tolerance will be answered, and I'm guaranteeing that you can think about the question. I'd love to see what the banking industry thinks about whether or not, for Morgan Chase, $13.1 billion is too much or too little credit risk, but that's an issue that we as an industry need to address. Can I take more or less risk? From an insurance organization point of view, we probably have more appetite for risk than we think we do, and it's because we haven't really understood how much we have across the organization. Of course, it ultimately leads to the question, can I use risk management as an offensive rather than as a defensive weapon? I'm certain "weapon" is not the right word, but that's the only word that came to mind.

Let's look at some core principles. Some of these relate to some of Uger's comments. There has to be a common measurement system. In my mind, having an independent calculation of credit risk and a methodology on credit, then combining it with an ALM or an equity or an interest rate risk measurement that's totally different and trying to pull the two together just won't get it done. Coming up with a common measurement system is important, as are looking at all the financial risks and eventually the non-financial risks, a holistic quantification of the overall financial risk exposure and then attribution back to the risk elements.

You want to define your risk elements by some statistical technique reflecting historical experience or your best expectations for the future. You need to recognize correlations across risks. What's interesting about the insurance business is that if you stop and think about it, how are the correlations? They're not highly correlated with each other. Even across products, across a risk element—I'll hit on that pretty hard—there is even diversification benefit. It's an important principle to capture into the whole measurement process.

What I'll do is make a lot of the points by walking through a case study. It's pretty simple, but it will make the point. To start off, we'll have three lines of business in our hypothetical company—universal life (UL), variable annuities (Vas) and a bank CD. This is an insurance company that somehow has a bank, and it sells CD products. The UL business will be backed by corporate bonds and the bank CD will be backed by mortgage pass-throughs, and that's by design. The VAs have some simplified death benefit guarantees but no living benefit guarantees.

The risk elements are interest rate, equity market risk associated with the VAs, credit risk associated with the assets, underwriting risk—mortality, in this case—and then lapse. What we mean by lapse in this example is that—we'll get into it a little later—there's an expectation for policyholder behavior. Lapse risk in this context is the risk that that expectation is not realized, and it turns out to be different from what you expected. We'll use a metric called earnings at risk (EaR), which we'll define in a minute, and we'll do it over horizons of one, five and 10 years. But as you look at this material, the results could be embedded value-based or any other measurement you might think of. It could be cash-flow based. You could even use a technique that's analogous to the C-3 methodology that's currently proposed.

To recognize all these risks, it's all built around the simulation approach for each of the different risk elements and the integration of those simulations. This will be analogous to your common scenario set. You're able to generate a set of interest rate scenarios, as well as equity return scenarios, using fully correlated scenario generation process. You can end up with a set of interest rate and equity scenarios that have a level of robustness and reflect correlation of those two elements. Then within that, you can also independently implement a credit simulation approach. In this case, what we've done is just used the coin flip from a distribution. We're taking a simple approach, and we can talk about some of the more sophisticated approaches you could use. Flip a coin. You look at the bond. It either defaults or doesn't, depending on where the random number put you in the distribution. You need to link the distribution so that you're not in the tail of the distribution for the AA credit and in another part of the distribution for the A credit. You want to make sure that you have some correlation across the different credit rating bands. You also want to worry about autocorrelation. In other words, if you have a default event in one year, there are probably some residual effects in subsequent years.

What we've done here is, again, a random process around mortality. We basically took the distribution so that the mortality rates reflect a long-term secular trend in which mortality worsens over time. For lapse, we're just looking at nonsystematic variation around the base assumption. What happens is that you end up with a scenario set, in theory, in which each scenario reflects an interest rate environment in the future, reflects an equity market environment in the future, reflects credit events across time in the future, reflects mortality events across time and lapse events. Each scenario is defined by those five elements. In putting these together, I mentioned the correlation between debt and equity. We've assumed credit events are uncorrelated with interest rates and equity markets, and surely that can be debated. Mortality is uncorrelated with any of the others. We made the assumption that lapse was as well. That doesn't mean that the base lapse assumption, for example, on any of the annuity products isn't linked to the environment. The likelihood that you mis-estimated on the assumptions is uncorrelated.

Then what we're doing in each of these simulations is projecting income expense and claims. These were done on a statutory basis, so we're producing book profit in each year over the one, five and 10. However many simulations you run, you end up with the distribution of results. You rank-order them from highest to lowest. Those of you who have done any simulation work will know that this is a pretty typical profile (Sabatini Slide 9). You might change this steepness of the curve, depending on the mix of business and some of the assumptions that you've made, but typically you'll find the earnings at risk is the distance between the expected values defined by the mean and the fifth percentile result. You can pick whatever point in the distribution that you'd like to measure. You could use the CTE measure for that matter.

Now we'll look at actual results (Sabatini Slide 10). This will be over a five-year period. We haven't done any adjustment. It's just the sum of the profits for each scenario. It's the sum of the profits over the five years. If there was $20 of profit per year, the value that would be shown for that scenario is $100. This reflects the distribution, and we want to look at just the Total column. This is the result across all risk elements. Basically the mean is $9.7, and if you take the $9.7 minus the fifth percentile value of $1.3, that gets us our earnings at risk exposure of $8.4. So, there's a 5 percent chance that I could lose $8.4 of my baseline earnings expectations over this five-year period. If you want, you can divide by five and think of it in terms of an annualized impact, and this is across all risk elements.

Now let's talk about each of the other elements. What we did in the Interest column is keep all the other assumptions constant. Mortality is no longer stochastic, credit is no longer stochastic and lapse is no longer stochastic. All we're doing is measuring the impact across interest rates, recognizing that we just have expected results across the other risk elements. We end up with a distribution. We get a mean and an EaR number. We do the same thing across mortality, lapse and credit. Now we'll do a little bit of audience participation.

Sabatini Slide 11 is a restructuring of those results. What it says is that you have—let's talk in millions of dollars—$16.8 million of total risk exposure, of which—it just turned out this way—$8.4 million is the diversification and correlation effect. If you measure the risks independently and add them up, you get $16.8 million. If you measure them in aggregate on an integrated and aggregated basis, you end up with $8.4 million, with the correlation effect being $8.4 million, so the correlation effect drops out of the calculation.

I'll offer some interesting observations. Over a five-year period, interest rate risk on a UL product is about 50 percent of the total exposure. That's not surprising. Because of our assumptions about deteriorating mortality, that's a much higher percentage than you might normally see if you just kept the table constant. Lapse is relatively small. Credit is relatively small. In particular, the credit piece is small because credit diversifies over time. Typically we don't have five years of bad credit markets. We have credit events followed by periods of relatively benign experience, unless you're pretty bad at buying bonds. There are techniques that allow you to take that correlated value and then attribute it back to the risk elements; I didn't do that here. There are some pretty precise statistical techniques that get pretty involved in doing that, or you can take some simplistic approaches to it.

Once you've done it across one product, let's do it across all the risk elements and the three products in our study (Sabatini Slide 12). The UL column is the result that I just showed you. In the variable annuity column, there's a ton of equity risk exposure, EaR exposure—a lot of earnings volatility. That's not surprising from VA business. In this example, because of the way we modeled the guarantees, what's really going on is the relationship of the fee income to the expense structure that's built into the simulation. There's a little bit of lapse, a little bit of mortality risk because of death benefits, and some interest rate risk because of the correlation between debt and equity that's assumed. A CD here is all interest rate risk. I can't even recall when we did this work whether or not we had any kind of lapse exposure on the CD.

Clearly on the VA business the correlation effect isn't as strong because it's being driven by that single risk driver, the equity markets, but there is some correlation. People don't always die when the guarantees are in the money, even though you might think otherwise, given some of the recent experience. So the VA business has less correlation on a percentage basis than does the UL business. Note that the correlation effect on the UL business is being driven very much by the fact that the mortality and the interest rate exposure aren't correlated at all.

If you go across products, I was surprised. I've seen more dramatic results, particularly if we had a single premium deferred annuity (SPDA) in here. Let's, for example, have an SPDA business. The scenarios that might drive some of your poor outcomes on the UL business are not necessarily the same scenarios that will drive the outcomes on your SPDA business. It depends on the guarantees, crediting strategies, how long the assets are relative to the liabilities and so forth. But typically what you find is that you'll get some real diversification benefit when you start looking across even general account products. It's not nearly as strong as you might see from some of the totally uncorrelated elements, such as mortality and interest rate risk, but it's stronger. There's a little bit here on interest rate on the VA and, of course, UL and the CD, and there's a little bit of correlation that you're getting at a lapse.

The key thing to note here is that the uncorrelated total across the risk elements and across the products is $42.7 million, but the correlated total is only $17.9 million. Now, this is an example, so we made a lot of assumptions. I won't suggest that the results will be this dramatic for your businesses. But in some of the live work that we've done, we've seen some significant benefit from diversification. It depends on your products and how much relative mortality risk you have and how your assets are positioned against some of your interest-sensitive liabilities. But the point is that at the end of the day, 25 percent of the risk exposure is interest rate, 50 percent is equity market and the balance is split between lapse, mortality and credit. Of course, credit is relatively small, again because of time diversification.

This is the distribution that's associated with the results on the prior page (Sabatini Slide 13). I just show this because it sets something up in the next slide, but the point to note is that we're measuring risk as earnings lost relative to the mean value, but you can't ignore the tail of the distribution. I want to make that point. We performed the same measurement but over a one-, five-, and 10-year horizon (Sabatini Slide 14). It allows you to start talking about risk and earnings volatility over a longer horizon. One of the things to take away from this is that credit is much more material in the short term. You would expect it to be. Equity risk is material, no matter what the time horizon is. Equity-driven risk is the tail of the distribution. It doesn't diversify over time all that much. Interest rate risk compounds a little bit. It allows you to have some very interesting discussions around how much risk you want in each of these risk elements.

Again, these are purely illustrative, but the idea of taking more credit risk and a longer-term view, somehow positioning yourself to deal with that one-year event or maybe pay some money to hedge out the one-year event but take the bet longer term and see what the net benefit is, makes for an interesting discussion. This would tell management that they have twice as much equity risk as they do interest rate risk. It would be interesting to see how a management group would respond to that. I would imagine that they probably intuitively understand that, but quantifying it really helps them.

The real value in all of this is using any measurement framework that you develop to do two things. One is to make decisions and to formulate some positions on how much risk. In this case (Sabatini Slide 15), we took the same case study and did a couple of things. The first column, the VA emphasis, is the results that I've shared so far. In the next one, I took those same products but changed the business mix and said, "What if our balance sheet looked dramatically different?" We significantly shifted the mix of business away from annuities toward UL, I think, and that changed the distribution. Then in the third example, I kept the mix the same, but we got more aggressive with our investment strategy around UL, and we extended the inherent mismatch. By moving to a balanced product strategy, the expected value or mean is increased dramatically. The tail of the distribution is more palatable than it was on a VA emphasis, and my risk is pretty much the same.

Of course, there are a lot of other things that factor into this. Could we ever get our balance sheet to look like this? Stop and look at some of the acquisitions that have been taking place over the years. Some of them have been driven intuitively by some of the thought process that's reflected here. In terms of the mismatch, you can see you get a fairly dramatic improvement in the expected value by just driving the UL business out to a more aggressive strategy. At least this risk metric says that on a per unit of earnings over a five-year period relative to the risk, we'll get more value by adapting this strategy.

There are many other things that will drive an ultimate decision. These are just two simple examples of how one could use some of this information to make some very interesting strategic and tactical decisions and, for that matter, in terms of doing some of the disclosures around risk. I think we'll get there. You've seen some organizations in their public disclosures start to disclose risk-type information. It will be a slow, evolutionary process, but I think we're headed there. Some of it's just the technology that allows us to do all of this. It takes a lot of effort. I wouldn't even begin to estimate what it would take for any reasonably sized insurance organization to get something like this done. This was a case study that allowed us to use some simplifying assumptions.

That leads into the implementation strategies. First of all, if you're going to do something like this, management has to want it. What I'm seeing in the industry is that it is starting to happen. Companies are starting to say that our approach to managing risk will be a differentiator in the marketplace. They're investing, and management is giving the support to the people who are needed to implement this. As for methodology, you need to work through the details. Make sure that everything is being done on a consistent basis. One of the things I like about what I've illustrated to you today is that it's the same methodology, particularly in terms of the ALM risk and the credit risk. The measurement approach is consistent, which is one of the challenges.

In an actual implementation, I'd probably mix and match on day one to get something implemented, but ultimately the long-term goal, to the extent that we can leverage the technology, is to get everything on a consistent basis. I've mentioned the modeling infrastructure. You can't integrate and aggregate risk and create those simulations. Reality says that when you start mixing interest rates, equities, credit and mortality, the number of scenarios begins to compound. You need to compromise waiting six months for the results against the need for precision. You'll have to make those kinds of tradeoffs, which is one of the advantages to using correlation matrices, because you don't have to compound the simulation process.

MR. GARY HATFIELD: One of the things you talked about as being critical was consistency of measurements. One of the issues that I think we're seeing in the VA world is inconsistency of measurements. You have guaranteed living benefits that are marked to market under Financial Accounting Standard (FAS) 133. You have other benefits that come under the Standard of Practice (SOP). You also have the base mortality and expense (M&E) fee income stream, which is handled under FAS 97. They are very inconsistent. Could you just speak to your perceptions on how much education is occurring and how companies are coming to deal with this?

MR. SABATINI: The accounting is the accounting. If I were a member of the FASB, we wouldn't have what we have today. But I am not. It's a difficult problem. It does carry over into the measurement as it relates here because if you decide on a statutory or a GAAP measurement system—there are some people who are looking at GAAP earnings volatility as a risk measurement—you need to bring all of these different accounting frameworks into the mix. The accounting is what the accounting is. If that's what will drive your GAAP earnings volatility and that's what you're measuring, you need to use the different bases. The fact that we don't like the inconsistency in the accounting systems means that maybe we should campaign to the Academy and the Society of Actuaries to get more involved with the accounting bodies to influence a more rational framework—one or the other.

MR. KOYLUOGLU: Perhaps I can add one thing. I think you need to define your economic view and stick to that, and think about all the accounting issues as constraints binding you in your evaluations. This is a very big problem in life insurance because of the long-term time horizon and all of these accounting rules. It's also a big problem in banking, but the banks prefer sticking to the economic reality. They see risk aggregation as a decision support tool. It doesn't really answer all the questions, but it answers a set of questions concerning the economic view. Leaving the accounting aside, in Frank's example we have seen different time horizons studied. You also need to think about what time horizon is the most suitable one for a given decision. What I'm trying to say is that the economic view that you use for decision support will be tailored to the corporate finance question in which you are interested. As the question changes, you might want to use 10-year time horizon versus one-year time horizon. You might want to get closer to the accounting reality versus the economic reality.

MR. ISADORE JERMYN: I have a question regarding operational risk. There wasn't very much discussion about how that's brought into this analysis. I realize that it's a lot more difficult, but I'd be interested in any comments you have in that regard.

MR. KOYLUOGLU: If you're thinking about enterprise-wide risks, it's part of the puzzle, and the industry has started to attack that lately. It all started in the banking industry, and now Basel 2 is asking banks to look at operational risk carefully and actually implement operational risk models. In any financial institution—be it life, property and casualty (P&C) or a bank—there will be limited data for operational risk assessment, so we definitely need to trust on the use of external data. The tricky bit for the actuarial analysis is the blending of internal data with external data for the quantification of your operational risks. You probably haven't experienced some of the big losses that some other insurers already have experienced, and there are lessons to be learned from those. That's great. But numerically taking all the data and customizing it to your institution is a blend of art and science.

In the past five years on the banking side, there has been a lot of studies and publications that you can take a look at and learn from. There are data vendors selling the data of events starting from the 1980s in the financial services industry world. Those databases would have almost 1,000 events. There will be a cutoff. All these events are greater than $1 million in losses. They would need to have a cutoff because these are all publicized events. There will be some sort of a bias in the underlying data because these are the only publicized ones; there are all the publicly undisclosed events, and you need to consider these in your calculations. There are a couple of really good studies, especially one led by Boston Fed on the banking side.

Once you quantify the operational risk distribution for your institution, the next thing is the aggregation. Usually people would assume very small correlation. Sometimes they won't really test the correlations. They will just increase the correlations because if there's an operational event and it will be publicized, there might be some effect on it on lapse rates, for example. You would like to consider these things in your calculations.

MR. SABATINI: I'll just add one comment, building on Uger's comments. Operational risk for the insurance industry probably can be boiled down to just two or three types of what I would call nonfinancial risks. The distribution and litigation and reputational risks are probably the big ones, as opposed to some of the other financial services industries. There's a lot more inherent operational risk that's a bigger issue that they can lose a lot of money on. It's different for the insurance industry, I think. I make all these comments with the stipulation that I'm not an expert at operational risk yet. It comes down to a frequency and severity issue, particularly in terms of the litigation issue. We have some data, but, of course, it's clearly not statistically significant. But if you can define your distribution, you can certainly then add it into the simulation process, decide to what extent it's correlated with anything else and factor it in.

FROM THE FLOOR: I have two questions. In the work that Chase and others are doing, are actuaries involved? Who's doing that analysis?

MR. KOYLUOGLU: As far as I know, there aren't that many actuaries involved. In the J.P. Morgan Chase calculations, this is often done by bankers with a background in economics.

FROM THE FLOOR: Interesting. My second question involves correlation, the whole thrust of your discussion this morning. Frank, I don't know if you said you weren't dealing with correlations in the tail or you were. If you make management decisions, it seems to me that you also have to be worried about correlations in the tail. You probably don't know what the sensitivities are or how they really work. As you pointed out, you can use a lot of Delphi techniques. You make your best guesses. That's probably what we're doing because we haven't had any experience and may never have any experience in these areas—and then to make decisions about products, about something you know a little.

I love the idea of using sensitivity. Sensitivity surely is useful, but they'll do sensitivity to find out what's most important. But how do you know what is right when we don't know how policy owners behave? There are obviously arguments that as you approach the tails, everything becomes independent. There are other arguments that you don't. What does correlation mean when you get out into the tail, and what does it mean to do sensitivity tests? How do you know that the results that you're getting have any use to management in making decisions?

MR. SABATINI: If you provide management more information and it's reasonably credible, it's better than anything else. Let's start with that. I'm sure you'll agree with that.

FROM THE FLOOR: Right, but bad information can get bad results, too.

MR. SABATINI: When you build your distributions, particularly around policyholder behavior, there's the whole contagion effect. You issue a guarantee. You have a diverse policyholder group. Not everybody will exhibit the same behavior at the same time. You're leveraging on some inertia. Then, some third party shows up and starts making a living out of getting these people to exercise, and you have 100 percent exercise. I think that that's part of the construct. As you build this methodology, being able to introduce that thought process and having regime-switching lapse behavior, having credit events in the tail and some of the correlations among different credit instruments—having all those correlations go to one is certainly plausible. From a technology point of view, it compounds the problem, but it's certainly viable. I agree that it should be reflected either in the base case or as a sensitivity.

MR. KOYLUOGLU: The uncertainty around dependencies is a fact of life. First of all, just by adding things up without any correlations, you at least establish sort of the worst-case scenario. The next thing is that it's your best guess for a point estimate. There are some institutions that are not really happy with that. In those institutions, the actuarial type of people would like to introduce uncertainty in the correlation matrix because there are some correlations you feel pretty comfortable with, but there are others that if you have a point estimate of 35 percent, you think that there is wide variation around that. In those institutions, they introduce this secondary level of uncertainty to the problem, and in their simulations they also simulate from a random correlation metric. So analytically, they just make the problem more complicated by bringing in that some correlations are pretty certain and some others are very fuzzy. So, let's introduce some secondary uncertainty around that.

In the end, this is a decision support tool. If you step back and look from the senior management's angle, most financial institutions don't have any tool yet, so they don't know this overall picture. This is a starting point. The results are fine, but I think the value is in the discussions, in the education of the senior management and in embedding this into policies and processes. The caveat is that you shouldn't forget the uncertainty around the calculations and correlation matrix and all the other parameters.

MR. WALTER NEEVES: My question is on correlation of market risk and credit risk. If you could, tell us what methods you've observed companies using to study correlation between market and credit risk and if there are any best practices in this area. Then additionally, if you could, please comment on the observed variability in the correlation between market and credit risk over time.

MR. KOYLUOGLU: Credit risk is a little tricky. The other tricky bit is the lapse, but let me leave that aside. What do we mean by credit risk? Is it just default, or is it default and credit migration? Is it default, credit migration and the spread volatility? To answer your question, I'll just focus on default, for example. In terms of market risk, again, let's just focus on interest rates and the equity market. There are two different approaches that I have seen embedded. One approach is coming from this component approach. You will run your different models for credit risk, for equity risk and for interest rate risk, and then bring them together. Bringing them together, people would define a proxy, and that proxy will resemble the core risk. For example, for default risk their proxy could be Moody's speculative grade of default rates. That will be their proxy. For equities, the proxy could be S&P 500. For interest rates, you identify that your portfolio is highly correlated to the parallel shift in the yield curve. That will be your proxy. Then you will look at the correlations across these proxies to come up with the correlation matrix. That's one way of doing it.

The other way of doing this is starting from the risk driver approach. In the risk driver approach, again, you could simplify it and come up with a set of market and credit risk indicators. Market risk indicators could be 10-year rates and monetary policy. The 10-year rate is the level. Monetary policy, which is 10-year minus the short rate, is the steepness of yield chart. That is defining your interest rates. Now for credit, again, if you are happy with an external benchmark, you could use Moody's speculative grade default rates. You will look at the correlations, generate a consistent set of simulations and apply these in your calculations.

These are two different approaches that I have seen put in place. Again, this is the starting point of the discussions. As people implement, they will say, "Moody's speculative grade default rates are fine because it's rich in data. But when I look at my portfolio, on average we are not really speculative grade. We are in between speculative grade and investment grade. So for investment grade, will I be able to find another proxy, and if so, what is that proxy? Is it really the migration, for example? If it is migration, then what is the right measure?" These will be the sorts of analysis that people go through to come up with the correlations.

MR. HARRY PANJER: I have a couple of comments, Frank. One of your slides presented different time horizons. The question is about time horizon. You said one year, three year and 10 year, and you presented the results. You showed that the capital at risk or the earnings at risk varied quite dramatically over those horizons. I tend to think that those kinds of presentations are somewhat misleading because they're done at fundamentally the same probability level for each horizon. In other words, you're comparing a 95 percent failure in one year with 95 percent failure within five years. It seems to me that the more sensible presentation is actually to comfortable with, and then look out to see what the corresponding probability is at the other horizons. So if I choose 99 percent for one year, I get this capital level. That tells me it's 97 percent at three years or 95 percent at five years. It's a much more meaningful exercise, I think, doing it that way. When you get these large variations in numbers and capital amounts required for different horizons, it then throws up the question of, what do I think the probability should be?

The second one was your comment about the question about the variability of earnings under two approaches, one essentially marking cash flows to market and the other one using essentially distributable earnings based on statutory reserving and so on. I think those are the two. You discussed those, and you said that you need to make things consistent. My question is, have you done any studies to validate that statement? It seems to me that if you do this in a closed block framework, you'll throw up all the earnings at some point. A statutory system will only determine the incidence of those earnings. The changes in present values may not be that much different, and I just wonder how different they are, based on your experience.

MR. SABATINI: On the time horizon question, I agree with you. It was a simple example. The AA bond default rate over a one-year horizon is much different from the AA bond rate over a 10-year horizon. I think that it's a matter of preference and the purpose of the calculation. But certainly looking at different points in the distribution over different horizons makes sense, which raises some of the issues around the whole C-3 methodology. If you extend the C-3 methodology to all products, one of the things that you'll observe is that the smallest present value will occur in different time horizons across different products. How do you decide what's the right CTE for what product? Harry, you raise a very good question. I agree with your point, and I think that it's one that we need to think about. I'm not so quick to look at bond default rates and somehow have that guide me as driving the right point in the distribution. I'm not sure what the right answer is, quite honestly, but it's certainly an issue that we have to address.

As it relates to metrics, there are any number of measures that you can use. In all honesty, when I do this sort of work, I typically will use embedded value as my primary measure, which has a much different horizon. That's a preference. I would tend to use earnings volatility as a constraint in the calculation. I don't want to put my value-at-risk (VAR) subject to some constraint around earnings volatility where the earnings volatility is a constraint in terms of evaluating particular options. I wasn't implying that a mark-to-market methodology, which is what I think you were asking about in terms of cash flow, is directly comparable to a statutory-based projection, although I think it will give—I've done enough work to have it suggest that it's going to give—sort of the same answers. I just find the mark-to-market methodology in an insurance context difficult to communicate and difficult to understand. But I think you can argue that if I do embedded value across thousands and thousands of scenarios across different risk elements and do the same thing on a mark-to-market basis on a one-year VAR, I'll get the same conclusion. I don't know if that answers your question or not. I haven't been able to prove it out algebraically. It probably can be done, but I just haven't sat down and tried yet.