2005 Valuation Actuary Symposium*

Orlando, Fla.
September 22–23, 2005

Session 44PD
Enterprise Risk Management

Moderator: Mary Ellen Luning

Panel: Mary Ellen Luning
David Ingram

Summary: Changes in the financial services industry have heightened the interest in enterprise risk management (ERM). Management is starting to develop new ways to examine financial risks holistically in an integrated, companywide model. Panelists discuss issues involved in taking an enterprisewide view of risk management. Topics include quantification of risk exposures, including insurable and financial risk; considerations in addressing nonfinancial risk; modeling risk on a companywide basis; and ensuring consistency in assumptions and methodologies of cross applications.

MS. MARY ELLEN LUNING: I'm with Ernst & Young. I see a lot of familiar faces out there, which is always good. We're here today presenting "Enterprise Risk Management," and I'm going to give an overview and some scare tactics on why you should care and why you should take time to look at this. Dave Ingram is then going to give us much more detailed review of what's going on in this space. He has been heavily involved with the Risk Management Task Force and is now with Standard & Poor's (S&P). He has designed the protocol for how it's going to look at this, so I think you'll all enjoy that very much.

As I said, I'm going to introduce what it is for people who might not know what the acronym means, but there's been a tremendous amount written about it, so I won't labor on that too much. Then I'm going to touch on a few reasons why it's a hot

NOTE: The chart(s) referred to in the text can be found at
http://handouts.soa.org/conted/cearchive/valact05/044bk.pdf

topic, why you've seen a lot of articles about it in literature, why we are here talking about it today and why it made this agenda.

What is ERM? It's explicitly defined corporate risk strategies. Those are a lot of words that basically mean that you need a strategy. If this happens, that happens. You need to write it down; you need to articulate it.

The second thing is robust risk infrastructure. What does that mean? That means you have to have the models and the communication tools in place to effectively implement your risk strategy. You have to know what your risk tolerances are for every kind of risk, not just for the ones actuaries are used to looking at for everything.

Next is oversight and risk governance. You have to be able to make sure it's going to happen. It's easy. We could all take a three-day off-site with a white board. We can write down all of our risks and all of our mitigating controls. If it doesn't happen, it's not effective, and it was a big waste of time. Governance rules are important to who's going to be responsible, who is going to get the reports and what you are going to do if it doesn't work.

What does it look like? ERM is a little different for actuaries to think about because we have to think about risks we're not used to thinking about. Credit risk, market risk and insurance risk are all the ones that we're used to thinking about. We have them on our to-do list; we have models; we have people who are assigned to those risks around the organization. Operational risk is a little bit different for us. We haven't thought much about it, but if you think about a lot of the issues that have gone on in the industry of late, a lot of it boils down to operational risks and not having the right kinds of approvals in place and oversight that you need.

What does operational ERM look like? That's a nice thing to say, and you read all the articles, but what are companies doing? Luning Slide 5 is sort of a picture of what that would look like for a particular organization. I'll walk you through it. Remarkably it looks a little bit like 404. I bet everybody here knows what 404 is and has had to suffer through some part of it at some point, right? It's a little bit like 404, only now you're covering all the risks of the organization, not just your financial reporting risks and your what-can-go-wrongs in financial reporting. You're covering what can go wrong everywhere. What can go wrong in pricing, what can go wrong in IT, what can go wrong in legal, what can go wrong outside of the company—you're basically covering everything.

The first thing to do is identify those key business risks, and then you want to identify the mitigating controls that you have to hopefully protect yourselves from those risks. You want to assess those controls just as you do for 404, but it's applying to everything else. You want to monitor the progress. Just as you have remediation on 404, you want to have an action plan wherever a particular risk is not mitigated to within your tolerance, and you want to monitor those action plans

and move toward your appropriate position. You want to continuously enhance and update your mitigating controls and your risk framework.

Why should you attend this session instead of cutting out early, having lunch at Universal Studios and then catching an earlier flight home before the rain? All things aside, if we had all the time in the world, ERM is good for business just by the way I just described it. It's good communication. It's good to know your risks. It's good to be able to communicate them to your employees, your boards and your stockholders. It's good business sense, but I would guess that most people don't have time for the things that are good for them. Like me, I don't have time for all the things that are good for me. I know that. How many people are busy? You have to have regulations step in sometimes to push you into doing something that you know is good for you anyway.

Basel II, as some people might be familiar with, is a banking rule and requires some management of operational risks for the banking community. In Australia, AFRA requires formal documentation of the risk management strategies across all types of risks. The U.K. Prudential Source Book is something I'm going to get into in a little bit more detail this morning. That is a capital requirement similar to the ones that we have, but in this case it gives you some capital relief if you have good risk management processes, systems and controls in place. I'm going to go into that in a little more detail so you can see what kinds of things are out there and what kinds of things might be in store for us at some point as they move across the ocean.

Section 303A is for us and has to do with corporate governance. It doesn't necessarily say that you have to have ERM, but there are some implications to it that could be interpreted that way. I'll cover that a little bit. Then, of course, the rating agencies have been looking at it, and so Dave is going to present a lot more information on that later after I'm done here.

Let me go over a little bit the U.K. Prudential Source Book. It's a requirement of the Financial Services Authority (FSA). It's required for all financial institutions, and, as I said, there are two areas of focus. There's the capital requirement, which is fairly similar to what we have here, but then there's this credit for systems and controls around risks. The other thing that's interesting about it that gets everybody's attention is that it holds individuals responsible, not just organizations, and that's fairly similar again to what we see in 404 where CEOs and CFOs have to sign an attestation that they have a well-controlled environment. Individuals are on the hook for complying with these rules.

The capital and solvency requirements require stress testing and scenario testing. In the United Kingdom this is going to require a lot more modeling and data capture protocols and things. They're not necessarily ready for that. I think a lot of people in this room over the past two days probably went to one of the sessions on variable annuity (VA) hedging or modeling and know from your own experiences

that that can be a big headache. We're probably a little bit further along in that, so I'm not going to go too far into that kind of detail. But those rules are not all that different from some of the rules we've talked about here at the NAIC and other places.

What is different about it is that it allows this relief for good risk management practices, and it says that these systems of controls around risk have to cover all the usual culprits: credit, market, operational, liquidity and insurance risks. Again, we're used to looking at a lot of those risks, but not so used to looking at the operational risks. In terms of operational risks, they're pretty clear about some of the things. These are just some examples, but there are others in the rule. Documentation—how many people who are busy have good documentation? You can imagine that's going to be a big effort to be able to document all these different processes.

Business continuity management is next. You could say that we should be more ready for that because of things like Katrina and other things. Most companies have some kind of business continuity strategy, and if they don't, they're probably behind the curve on that. Employee responsibilities are a new one. If you think about all the different issues that have come up and things that you have read in the paper over the past several years, almost all of them can be tracked back to a person who either took an action or did not take an action to cause whatever the issue was. Controlling people is difficult, but controlling the actions of your employees and having the right oversight, the right sign-offs, the right protocols, the right reviews and producing that information are a tremendous effort.

Outsourcing is something we do to save us time, but there's a cost to outsourcing, as well. You have to figure out how you test the controls of that firm that you've outsourced to. Information technology (IT) is another area where I think we should be a little bit ahead of the curve because there should be good controls around that, although I can imagine if you looked at anybody's IT shop, you'd find some improvements that could be made.

External events are an area where you might do scenario testing. The whole world has been criticized for not predicting what would happen when Katrina hit, but that's an area where if you build models and test them and then act on what you see, it could be helpful. Information security and privacy rules are something that we do deal with in the insurance industry often, and again it should be an area where we should have some good controls already in place.

The other thing that the Source Book requires is a written risk policy, and it's required for every risk. You have to have all your details in there: what the risks are, exactly what they are, who is responsible for controlling them and what the mitigating controls are. In addition to that, you have to talk about how your internal assessment process works and why you think it's effective. Why is the honor

system good in this case? What do you do in terms of compliance and governance to make sure that people are obeying these rules and following your strategy?

The other interesting thing is you have to show exactly how your capital models are linked and how they're used in your strategic decisionmaking. Are they used for allocating capital in terms of performance measurement? Are they used for individual management or performance measurement? CEOs get their compensation based on capital models, so it's important. If you don't have that link, the relief that you could potentially get would not be granted.

In ending on the Source Book note, once it starts in the United Kingdom, you never know when it's going to jump on the Queen Mary and come over here, so we should be ready for those kinds of rules ourselves.

Section 303A is governance rules for New York Stock Exchange-listed companies. It lays out the requirements for audit committees and for board members, CEOs and others. One of the many things it says the audit committee of the board should do is discuss policies with respect to risk assessment and risk management, and that's all it says. It doesn't say that you necessarily need to have ERM, but it says that the audit committee is responsible to discuss that. If I was on an audit committee and was responsible to discuss that, I'd want to have the state-of-the-art ERM process in place to be able to protect myself. It also says that the compensation committee of the board needs to review and approve corporate goals and the CEO compensation. Again, some people have speculated that these rules will lead audit committees to require an ERM-type platform and lead CEOs to want an ERM-type platform if their compensation is going to start reflecting their controls on these other risks they're not used to dealing with.

The other cynical comment I would make is the first time somebody sues a company because it didn't have a good ERM framework in place, it will get even more attention.

Next I'll talk about some scare tactics to get people riled up about ERM. I want to talk a little bit about what the industry is doing, and good old-fashioned peer pressure is always a good way to get people involved in something. Ernst & Young did a survey of asset-management firms. These asset-management firms, especially the ones in Europe and the ones that are associated with banks, have had to implement these things because of regulations. But the survey came back, and 79 percent have a chief risk officer (CRO); 75 percent have a centralized risk management group; 60 percent have a risk governance committee; and a whopping 74 percent either have an enterprisewide operational risk framework or are in the process of implementing one.

We're in the process of doing a similar survey for the insurance industry to see what's going on out there, but I suspect that a lot of companies are going down this road, as well. In fact, I'm going to take a quick survey right here in this room

because we have a lot of people, and I won't have to do the work for the real survey. Does your company have a CRO? There's a lot. Do you have a centralized risk management group? How about a risk governance committee? It looks like a couple. The last one I think is going to be a lower show of hands, but how many have an enterprisewide operational risk framework similar to what I described before? We have quite a few.

MR. INGRAM: As we're getting the microphones ready, I want to take a poll, as well. How many think that talking to a rating agency is more fun than going to the dentist? We have a couple; that's good. I have a friendly crowd then. We are working on a process that will make that fun last a little bit longer, but hopefully as your dentist might tell you if he comes up with a new process that will extend your time with him, this process will be something that may make some of your visits longer, but may mean that you'll have fewer of those root-canal-type visits. We're all in favor of less root canal.

S&P started about a year ago working toward a process of incorporating ERM into our insurance company ratings. We saw some good reasons for doing this, and many of those reasons align well with the reasons that insurance companies might consider doing this. After all, our basic objective is to get our view of the financial strength of an insurance company correct, and we feel that the way that ERM approaches that, the holistic approach to it, the advanced type of modeling and analysis that is sometimes incorporated with ERM, that all of that will help us to do a better job in our analysis of companies.

Ultimately one of the things that we expect to do from this process will be to get to the point where we will have a better understanding of the capital needs of companies. We are going at this as kind of the opposite point of view that Mary Ellen just described that the FSA in the United Kingdom is doing where its focus is primarily initially on the capital models and tying ERM into as coming in through that door. We're going in the opposite direction. We think ERM is important and see looking at the capital models as a by-product of that work. Our objective in this is to make our rating process better.

Right now we have a rating process where we look at seven major factors in a company: competitive position, management and corporate strategy, operating performance, capitalization, liquidity, investments and financial flexibility. Within each of those components we do look at the risks and the risk management of the company now. Those blue bars at the end of each one of those bars (Ingram Slide 4) represent the idea that risk management is throughout our process now. Let me comment that the length of these bars was meant to illustrate the fact that for any one company we will view each of these things as having different strengths for that company.

What we're doing with this ERM process is we're taking these risk management things that we've always been looking at and gathering them all together. When

we've done that, just as a company does when it starts doing an ERM process, it finds that there are gaps. What we hope to do with this process then is fill in those gaps, and that is exactly the same kind of thing you do in an insurance company ERM process.

There's a good bit of overlap between what I will be saying and what Mary Ellen said, but I think I need to do that to give you a full picture of where we're going with this. What do we think good ERM is like? If you thought of this as a chart of all the risk management that was going on in a company (and I have to admit I borrowed this from somebody else), one thing you can do is draw a circle around it all (Ingram Slide 7). We don't think that's good ERM, just taking existing things, drawing a circle and saying, "Now I have it; we've done it all." We see ERM as more a process of organizing all those processes in a holistic consistent way. Ingram Slide 8 is an example of that. It comes from our banking ERM process. It's not a suggestion of what an insurance company should do. We are doing the same thing in the banking industry, and while from a high level our processes look a little different, when you get down into the details, the details of our processes and our criteria are almost identical.

What is this ERM? Why do we stick that word "enterprise" on there? What's the difference between that and just risk management? First, let's define risk management. See risk management where you're putting together a control process; you're monitoring, limiting and managing risks with the idea that you want to limit losses. The product of risk management, if you do it right, is a controlled environment of risk taking, and that graph in the corner (Ingram Slide 9) is meant to show a picture of a controlled risk-taking environment, where you haven't eliminated volatility. That line that bounces around is meant to represent the volatility, but the risk management is those lines above and below that that are showing the boundaries on the volatility that you are trying to develop through your risk management program.

We see there are two characteristics of ERM that go beyond risk management. One of them is that the ERM goes throughout the organization, throughout all the risks of the organization and consistently across all those risks and is also done in a way that's tied into and consistent with the fundamental objectives of the enterprise. ERM will not look the same in any two companies. There is not a template of what ERM is. That's one of the things that makes the discussions of it difficult to do in a broad way, because your ERM is going to be different from another company's, even if that company has a business profile similar to yours. The ERM needs to be consistent not only with the businesses and the risks that you take, but also with the organizational structure, the governance structure of the company and most every other aspect.

The second thing that sets ERM off from just risk management is the upside part of it. ERM, when it's fully realized within a company, has the upside of providing companies with a language and a basis for making decisions that will allow it to

ultimately optimize its risk-return profile and to use the information and the language from risk management in strategic decision-making, in capital budgeting, in product design and pricing, in investment selection and also in performance evaluation to work on improving the company's risk-adjusted return.

S&P will be publishing and implementing soon a process for evaluating ERM as part of our process, and that process when it is implemented will apply to all insurance companies globally, all types of insurance companies. We'll be implementing it starting a little bit later this year. If you don't hear from us about it in detail this year, you're on the schedule for next year. It is coming soon. What I want to tell you now is a little bit more of the details of it. Within the next month we're expecting to publish our detailed criteria on this, and our objective is to be totally transparent in how we're doing this so that if there's some way in which you feel we aren't explaining what we're doing, we are looking for questions and responses to our process with the idea that we will keep improving our disclosure of that to the point where it is totally clear to everybody what we're doing.

Our view of ERM, to be able to organize our process, is we've identified five major areas that we're going to be looking at. We do see the area that we're going to call risk management culture as being the basis on which an effective risk management process needs to be built, and the strategic risk management is the capstone of the process. One thing that is confusing in this picture is that the pillars in the middle of that process (Ingram Slide 12) are all the same size. It makes a nice picture but does not represent the idea that we think that they're exactly the same size, so as I talk about this, I hopefully will become a little clearer, but let me talk about the different pieces.

Risk management culture, the underlying piece, is the way in which the company has incorporated risk management into company decisionmaking. The comparison that I like to draw that seems for some people to make this idea click a little bit is you go through a process with the company where the company has increased its awareness of expenses to a great degree. You might be told that your spending will be monitored all the time, that it is important for all the people who have spending authority to pay attention to what they're spending money on and to be able to pay attention to the budget that they will have and to answer to their budget on a regular basis. They have an expectation that if they go over budget, they will have to have a discussion with somebody. You go enough over budget, and somebody gets to be further and further up the organization. You get the impression that it's important to the organization when you consider an expenditure that you make sure that you're getting the most for your money.

A risk management culture takes the same ideas and applies them to the idea of risk. Within a risk management culture throughout the organization, people who have authority to take risk on the company's behalf understand that they will have a budget for that risk, that somebody will be watching them if they go over that budget; they have to explain that. They will be looking for opportunities to be able

to make sure that when they take a risk, the company is getting the best return that it can for that risk.

We also incorporate into the idea of risk management culture how risk management itself is organized within a company. One of the questions that we frequently get asked when we talk about this is, What's right in a structure? Do we want to have a large corporate staff, or do we want to have a staff in risk management that is dispersed in the business units? Our answer is that we think a company probably needs some of both, but the exact mix of it depends on the company. Even more important is that the company has to be aware that either way of concentrating has its strengths and weaknesses. It has to have a way of thinking about how it's overcoming the weaknesses of the structure that it chooses.

Governance also fits under this—how risk management is reflected in the decisionmaking process of the company, how effective the risk management point of view is, and when it is incorporated into the decisionmaking process. Another major area that's incorporated is the idea of disclosure, both internal and external. The excellent practitioners in risk management are transparent about it. People within the organization know what the company is doing on risk management; they know what the risk positions are of the company; and they know what the company's intentions are on risk, risk limits and risk tolerance. Some of those companies are also doing a high degree of external disclosure on that.

This is an example of a couple of the criteria. As I said, we're going to publish something within the next month. This is a sampling of those criteria, things we're going to look for. Primarily in risk management culture we're looking for that commitment to risk management, looking at things like the quality of staff and communication, particularly with upper management and the board, looking for policies and procedures that are clear and well-known. Management compensation is another key item. There's a spectrum of how management compensation and risk management can align. In some organizations the management compensation incentives work directly against risk management objectives. In other organizations they work completely aligned with it, and in others they're neutral. Clearly, when we're looking at risk management, we want to see incentives that are at least neutral.

Finally, the last item is risk management monitoring independent from risk-taking and management. That is something that from our experience has not been a particularly high value in many insurance companies, and it's something that is an extremely high value within the banking industry. I've only talked to a small percentage of insurance companies that have paid attention to this. It's something that ought to be considered more carefully. The idea there is whether or not you need to have somebody doing the measurement of risk that's separate from those who are taking the risk. Is the person who's running your asset-liability modeling (ALM) program the person who's telling you how good a job he's doing? Somebody chuckled, but how many companies have a separate person doing the reports on

the ALM program from the person who's running it? Very few that I know of, and maybe as we talk to companies we'll find that there are a lot more. I don't know of examples in insurance companies where that has led to serious problems, but there are certainly a couple of extreme examples in banking where that kind of thing has led to a problem.

The biggest area as far as its scope within ERM is the risk control processes. There we're looking for control cycles. We're looking for companies that have identified, evaluated, quantified, monitored, diversified, limited risk, exploited risk and transferred risk, so on and so forth and also within this we're looking for how a company reflects risk in it's new product development. That's particularly important when the new product stage is one of the big gatekeeper stages for risk management that a company can drastically change its risk profile for the worse in a short time by introducing a new product that has been ill-conceived insofar as its approach to risk and risk management. A good practice company there will be one that, when it rolls out a new product, has the applications worked out, the policy forms, the administrative systems, the sales material and the approach to how it's going to manage the risk for that product in place the day it starts selling it.

Next we'll look at where this risk control idea applies. Ingram Slide 18 is a broad list of areas that we're looking at. We don't think of it as being *the* list. Every company probably has its own list this way, but this is the list we've developed to help us make sure that we think of all the things. The highlighted things—credit, interest rate, equity, insurance, operational—are probably where we're going to start in our first pass on this process talking to most life insurers. What we're expecting to do with this process is to tailor it to each individual company. The chart I have up there is identical to the chart that's being used by Office of the Superintendent of Financial Institutions (OSFI) in Canada and also identical to the chart that's being used by the NAIC in its new risk-based surveillance project. The idea that we have copied from them is that we want to concentrate our time, efforts and energy in looking at the areas where the company has a large amount of risk and where we suspect that their risk controls might be of lower quality. In our process, through the information we already have and through new information we'd solicit from a company, we'd form a view of the size of risk, what the most important risks of the company are and from our previous discussions we'd form a view as to whether or not there are excellent, very good or poor risk control process already in existence. What we will notice in there is "don't know" will be the answer on quality for some processes. From that we will develop our agenda. That will be an interactive process with each company, but the idea, as I said originally, is to tailor this to the company. Certainly in this process as well, we don't mean necessarily to exclude talking to a company about the things that it's doing well. We're not meaning this to be totally negatively focused, so in that process we do want to make sure we're capturing all the pieces there, but we're trying to manage both company time and our time most effectively here.

The criterion we're going to be looking for is quality—quality of the risk identification, quality of the monitoring process, standards and limits, enforcement of limits and effectiveness in execution of the risk management program. This presentation is being given this month in Latin America and Europe, as well as here, and one of the other speakers questioned what the last sentence means—effectiveness and execution? Effectiveness means that you have a control process that if you did it, it would work to control your risks. Execution means that we've found evidence that you are doing it. As we go forward in time, we will look carefully for interruption in execution. We will look for signs of that interruption because that will usually either be a signal that bad things have already happened, and so the risk management system that looked good when times were good will suddenly be something that you're changing drastically. That's one of the ways in which risk management will help us to avoid these root-canal-type things: by having it be a way of signaling problems in advance so that we can start talking to management about them before they become as serious and possibly before bad things become rating events.

I'll make one more comment on the risk limits. My story on this is that I personally, and S&P's view on this will be forming, think of the three bears standard for risk limits—a not too hot, not too cold, just right kind of idea. Many companies have risk limits in place, and the company has never gone close to hitting any of the limits. You can say that that's evidence of good risk management, or you could say that that's evidence that the risk limit doesn't have anything to do with their operations because a good limit system will be part of the way in which you manage the risk, and so if you're never hitting the limit, it probably didn't help you at all, you need another thing, and maybe you'll call that other thing a porcupine.

You never hit your limit, but maybe you hit your porcupine, and that tells you you ought to pay more attention to the risk. That porcupine ought to be placed close to where you're operating and at a point where you have some concern. One CRO told me that the company has a duration-matching limit and it's never exceeded it. I said, "Do you mean that the person who is the boss or the person who is running your ALM program has never talked to them about their risk position? If I were their boss, I'd have a sensitivity that if 0.5 years were my limit, if you were at 0.4, I'd want to talk to you. That's your porcupine. You don't have to call it a limit, but it's a point where you're concerned and start doing something. The idea of limits—there are hard limits, brick wall limits and other limits where you do some actual management, and what we're looking for is the management process in there.

Extreme risk management is another aspect of risk control, but it refers to the types of risks that don't fit under a control process. Those risks are low-frequency events, and so the fundamental part of a control process is a monitoring report. If you put in a monitoring report for low-frequency event, what feedback do you get? This period, didn't happen; this period, didn't happen; this period, didn't happen; this period, didn't happen. It's not useful feedback, so you need other processes to deal with your extreme events. Some of those processes (things like trend analysis,

stress testing, contingency planning, problem postmortem, risk transfer, and I'll just mention the phrase "process reporting," which shows up on a number of these slides) refer to the idea that one of the signs of good risk management is that you review your risk management process periodically, so that's kind of an icing-on-the-cake aspect of risk management.

The kinds of things we're looking for is people looking at extreme risk scenarios. The list on Ingram Slide 23 is not meant to be an exhaustive list. It's just meant to be a bunch of examples: economic, physical, environmental and man-made disasters. We caveat that within our view of this, we're not trying to scoop in a lot of other things that are usually somewhere else. Generally with an operational risk, as was mentioned earlier, you have things like business continuity considered a standard part of operational risk. We still think you ought to have business continuity work, and we're not trying to move that around on anyone's lexicon. Also, we don't mean that if your core business is catastrophe reinsurance, you ought to have this one-off process for managing risk. If your prime business is catastrophe reinsurance, you need to have a day-to-day process for managing your risk. The companies that survive in that business do have those processes. You need to have some ideas and some scenarios. You need to test those scenarios against the company.

We think of liquidity risk analysis as being one of the kinds of things that fit under that category and a lot of other stress testing. The example with a liquidity crisis that I've been told is if your company gets into a situation where it has a liquidity crisis and you start immediately to figure out what to do, the company will be gone by the time you start executing. If you think that liquidity is an issue that you need to be prepared for, you need to have gone through the drill of figuring out exactly whose job it is to do something and given them some clear guidelines of what to do and maybe not the exact list of securities they should sell, but maybe yes, the exact list of securities they need to sell and also giving them the authority to do that in an extreme situation.

One of the obvious examples everybody throws around is that of the General American situation. One of the things that happens to a company that is in a severe situation is that the kind people on Wall Street notice that you're in a disaster situation, and your personal spread suddenly gets much bigger than the market spread. That is something you don't notice from models, so that's one of the reasons why you need to have this situation in place because if you're out there raising liquidity before Wall Street figures it out, you may be able to execute the trades at an economic basis and get that money that you need.

Reputation risk planning is the same kind of thing. The classic example everybody refers to on reputation risk management is the old Procter & Gamble story with Tylenol; anybody who's taken a business class has heard that one. That's an example of a company that went through a crisis and managed the public perception of it in such a way that there was almost no blip in the company results.

There are many other companies with examples you can point to where management did not know what to say immediately and either said nothing and let somebody else frame the picture or said the wrong thing, which led to the company's reputation spiraling out of control. For financial institutions if you think you have another asset that's bigger than your reputation, I'd love to hear what it is.

Looking at crisis response rehearsal, an example of that is what's happened in the past couple of weeks with Katrina. We sent out a survey to companies last week asking them what their expected loss was from Katrina. The first company back was back in a couple of hours. That company wasn't a P&C company, so it had an easier job to do. It was a life company, and so it had no direct claims risk from that, but it gave us a six-page report detailing every investment it had in that area, detailing every supplier it had in that area, summarizing all the policyholders it had in that area, and four or five other ways in which you could look at how it was exposed. I don't think that company did that starting after the storm. I think that company must have had that idea in mind to do that before then and a procedure in place to be able to think that broadly that quickly.

That comes down to execution. We can't have as part of our annual review asking you how you would execute in a disaster situation, but what we will do when disaster situations happen is we will collect information on how companies are executing. The ones that do the execution like the company I mentioned with the immediate awareness of what its position was show a clear sign of having this process in place.

The next step in the review criteria is a learning process. We've already met with some companies' managements that were in the property business in the Gulf area. The best ones come to us with, "Here's what happened to us, here's what the result is, and here's what we learned from this." They come up with a list. Here are the five things we're going to incorporate into our forward-looking procedures that we've learned already in this, and we're making big changes here and here, small changes here, and we have a study we're going to do here, and that is a part of this extreme risk process.

Finally, if you have a good environmental scanning process in place, you'll be in the position of the economist (this is an old joke) that predicted nine out of the last five recessions. That's what a good environmental scanning process will do. The business judgment that has to go into that is, How many of those nine times do you get prepared for the problem that your environmental scanning process tells you is coming? I don't know whether you saw any stories about how Wal-Mart responded to Katrina. Wal-Mart started its response when Katrina was upgraded from a tropical depression to a tropical storm, before it was upgraded to a hurricane. It gives you the hint that maybe it started responding for hurricanes that never hit, but it is an example of a company that came through that process in extremely good shape. It was doing a better relief effort than the government by far from its

infrastructure and from its preparedness, so it obviously had made a business decision that it was going to respond to seven or eight of those nine signals that it got, even knowing that some of them would be wrong.

A lot of people associate ERM directly with models and complicated calculations. You can see from everything I've said so far our focus at this point, a lot of it, is on process, on management. Models are important, though, to this. Some of the things that we are doing in the insurance industry are sophisticated risks that you can't get a handle on without models.

At this point we're going to be talking more to companies. We already do some talking about this, but talking more about the models, making sure we understand what kinds of methodologies the company uses, how the models deal with risk mitigation, how the models deal with risk dependencies and aggregation, what different risk measures are used, how the assumptions are formed and updated, how the data is fed in, how the company assures the integrity of the model and how it validates and documents it.

Think of these measures as falling into two broad categories. For lack of a better name, I called some of them primary and some secondary. The primary ones I'm thinking of are things where you're trying to directly measure what risk is, and often that will be something like a value at risk (VAR) or a conditional tail expectation (CTE) calculation. We're looking for companies that have measures like that that go comprehensively across the company, across the enterprise, and the best companies will be moving toward having consistent measures that can be used across the enterprise. We're also looking for measures that are executed in such a way that the measurement is delivered both in a timing and a form that's actionable and that the company does act on, particularly thinking that risk management doesn't take place by clearly focusing on one measure, and that the risk of our products, at least, is complicated.

If you look at one point in time, if you look at one spot in the probability space, you are not capturing the risks that you have—you're capturing one aspect of those risks, and the total picture of those risks is at least three-dimensional if not more. So a point measure is by definition almost going to be inadequate. We are expecting, and we've seen from talking to companies, that the best company practitioners have multiple measures. Some of those secondary risk measures are the ones that are used more on a day-to-day basis to manage risks. I didn't put it on the slide, but they're the obvious ones—duration, convexity, delta-gamma-vega, etc.—the things that you use to push risk around some.

If you ever tried to manage your risk by using a VAR, a comprehensive risk measure, it's just as strange an idea as the idea that you could buy clothing based on a single number for your size. Are there any women that have tried to do that? Have you seen it not work real well? The idea is fundamentally flawed that you can do that. You have too much risk. What do you change? The number itself doesn't

give you any hint, so you have to do something else to do that, and usually that something else is some kind of shocking process to the system looking at deltas using the word broadly rather than specifically.

As I mentioned already, we look at adequacy of modeling infrastructure. My comment during aggregation of risk is that S&P's position traditionally has been that we do not give credit for diversification in our view of company's capital. We know that's not exactly accurate, but we haven't had a better view. What we are expecting to do with this process is to do a lot of listening, to hear from companies about how they're looking at aggregation, and see whether from that we can learn enough to form a view that we think we can go forward with that's less out of synch with reality than the view that there is no benefit of diversification. Clearly diversification is the most powerful idea in risk management, and we do recognize diversification effects within risk categories, so we do hope to get a position where we can do a more realistic job of recognizing that between risk categories.

Strategic risk management is the last area of our process. This aligns directly with the second half of my definition of risk management. Strategic risk management is the way in which companies are using risk and risk capital ideas to inform their whole decisionmaking process. I'll mention that, in talking to a number of companies, what I've found, for instance, is that a number of companies don't consider this to be part of risk management, but it is a part of management. We will probably react to it within our risk management part of our discussion regardless. On the other hand, a number of companies in their risk management structures have included what they call strategic risk management, which is the risk that their strategy may fail. That is something we think is extremely important and is included in our review of a company, but it is not going to be included in our ERM deal.

When we're all done, we will look at these five areas of risk management and form a view of the strength of each of these five areas. We will, with that process, then go forward. This mirrors exactly the kind of process we do across the entire company now, and we will then boil this down into a view of the strength of the ERM and the quality of the ERM, which will then become one of eight things that we're looking at in rating the insurance company. It will be part of our discussions with management. It will continue to be a part of our discussions and our rating committees and will now be an explicit part of our ratings reports of companies.

I mentioned the time frame on this a little bit earlier on. With this kind of discussion we're drawing toward the end of our exposure period in what we're doing with this. I've done this presentation privately with a number of small groups, and we're doing this in a number of places around the world right now. Within the next couple of weeks, we hope to take any feedback we get, incorporate it into our criteria piece and publish that. Our target is to have it published before the end of next month. Shortly thereafter we'll start implementing this. It won't be like a curtain going up. In the next couple of months we'll be thinking it through with each

company we're meeting with. It will be more likely companies where we're starting the process of the annual review, after the publication of the criteria we'll consider putting this into the process, but certainly by a year from now we will have incorporated it into the annual review process with every company.

We're planning during the first six months after we roll this out to pay particular attention to feedback we get on this and having a broad review of what we're doing on this and evaluating whether we're going to need to make any significant adjustments to it based on that feedback.

During that time as well we'll be working on expanding these criteria, getting more detailed in what we're looking at, because we recognize that what we'll be coming out with at this point in time is too broad brush for us to form an adequate opinion of the biggest, most complex companies. So where we're heading with this is to develop more detailed criteria that will be applied to the largest companies. We're thinking initially that would apply worldwide to 30 or 40 companies; we'll see as things go on where we think that our ERM review will take on a significantly larger portion of our total time with those companies.

Ultimately with the companies that we have done that more detailed review with that have economic capital models, we'll start exploring the idea of doing an additional process with them of evaluating their economic capital models in detail at a level that goes way beyond things that we have done before, because what we're hoping to be able to do is take the company's internal model results and incorporate those into our view of the capital adequacy of the company. To do that we feel we want a company first to have demonstrated good ERM practices, just as the Basel requirements are for that, and I think the FSA requirements in the United Kingdom, and then we need to be comfortable in detail with the model because we'll be relying on the that number directly.

MS. LUNING: I have a quick question. David, do you want to talk a little bit about what you mentioned the NAIC is doing?

MR. INGRAM: The NAIC has passed a set of model guidance. I don't know the exact time frame when that happened, but it's in the last six months or a year. The set is a framework for the states' auditing teams or whatever they're called, the teams that do the quadrennial reviews, to structure their process of reviewing. It's called their risk-based surveillance system or something. Under that system, rather than repeating what the auditors do and looking at the details of how you put together your financial statement, they will do this risk assessment of a company and use that to decide what it is that they will spend their time on in their quadrennial review and then focus in on particular areas.

Three states have already adopted this process. If you're in one of those three states, you'll see this soon and, in fact, I know that at least two companies have already gone through the process. Those three states are New York, Ohio and

Minnesota. Minnesota is the one that implemented this first, and as I said it's already done two companies in the review process. One of the reasons I am able to give this information is that I was talking to someone from one of the companies in the hallway beforehand who said the result of this process for the company, and it's a company that does have a highly evolved ERM structure, was that its review process with the state was much shorter and much cheaper than it had been previously. The state auditors felt better about the process when they were done than they had with their longer, more intense, more diffuse process. While the trend in the United States isn't to require the risk management by the regulators, it will be that if you're not doing risk management, you will feel more pain when you visit that dentist.

MR. DANIEL J. KUNESH: This is for David. What you discussed seemed to be highly qualitative measures and analysis procedures to qualitatively evaluate a company's ERM practices, but you said little about a movement towards perhaps internal quantitative benchmarks or guidelines that you're going to use, I think you're saying that perhaps you will evaluate the results of a company's economic risk capital and the quantified risks in that regard and make some judgments. That seems a bit nebulous from an evaluative standpoint, particularly when you look at companies in different countries and so forth. Could you comment on that? Are you moving in that direction? Also, are your competitor rating agencies moving in the same direction that you are?

MR. INGRAM: Let me take the last question first. I'd rather not speak for my competitors, and I think most of you would have the same reaction. You'll have to ask them what they're doing.

With regard to quantitative aspects of risk management, we do not expect to be developing more quantitative tests of risk and risk management because, particularly in the areas where the risks are the most complicated, we don't see an outside body as being able to have access to the right information to do that adequately, so we do not intend on coming up with our own economic capital model or any other risk model that way to judge that. Our best response is this process I described where we would expect to go through this entire process and in the end come up with a review of the economic capital model. If there are things that aren't economic capital models that companies want us to look at, we'll certainly talk about that. We're not thinking of this as the only way we will go with it, but this is kind of a theme that we have.

One of the reasons why we think that this qualitative way of going about it is important is because our ratings are primarily meant to be forward-looking. We're not giving a report card on how management did last year; we're looking to see what its financial strength will be next year. If you have the absolute best measurement of what risk was last December 31, that isn't particularly helpful to knowing what it's going to be next year unless you think there's a process in place that's going to give you some ability to predict what the risks are the company will

be taking during the coming year. An ERM process is exactly that, so that's why our focus is in the order in which it's going and the relative types of emphasis.

FRANKLIN C. CLAPPER, JR.: I have a comment more than a question, but maybe you would want to comment back on it. I'm supposed to be a valuation actuary, but for the past three years or so I spent the majority of my time on Sarbanes-Oxley stuff. I was thinking about this. I think Sarbanes-Oxley 404 is working against better execution of ERM for the simple fact that management right now is distracted by this. It's a high priority, and it's excluding consideration of other things that they don't have to sign on the dotted line, whereas with Sarbanes-Oxley 404, they have to sign that everything is under control, and so it's getting a lot of emphasis right now. I'm all for getting better overall risk management in place so we can have better balance going forward. I don't know whether you wanted to comment about that.

MS. LUNING: I think being a practical person, I've been there with you, Frank, and I think the development cost is high. There's a lot to do to first document your processes. It's much more than anybody believed it would be, I suppose. Once you have it done, I think it is good. It gives me more comfort that things are being done correctly, that the controls are there and not just a handshake kind of thing, but I also think that you have to know when to stop overthinking it. You get to a point where you have the good controls in place and have it all documented, and you're done. I think there's probably some extremeness to it, and I feel your pain, but I think getting through the development process is a good thing to do.

MR. INGRAM: I'd react that I think it's favorable to the development of ERM. Not that people have been distracted by 404, but that ERM hasn't been led in this country by regulators and hasn't been led by the laws like Sarbanes-Oxley. We could develop what would be an effective principles-based system to management risk and that maybe when it's well-developed, the regulators will recognize it and pay a lot more attention to it. But I think it's best for us if it goes in that order, so it's unfortunate that so much time has been spent on process where the processes aren't uniformly valuable to the company, but it's lucky that something that is important has managed to escape that for now.