

RECORD Volume 30, No. 3*

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

Session 60TS Enterprise Risk Management Tools and Analytics

Track: Risk Management

Moderator: Hubert B. Mueller

Instructors: Samir Shah
Fred Tavan

Summary: There has been a wide range of approaches used to measure, monitor and manage risk. This session reviews both the quantitative and qualitative tools and analytics companies are using when they implement enterprise risk management (ERM).

MR. HUBERT B. MUELLER: I think the market is probably a lot more focused on ERM than ever before. The recent inquiries into contingent commissions on the property and casualty (P&C) side and for group business are proof that there are lots of non-financial risks that we are faced with as insurance companies. Not all of them are easily modeled using stochastic modeling, unfortunately.

Our panelists will cover both financial and operational risks. Fred Tavan an assistant vice president at Canada Life in Toronto. He has more than 10 years of experience in the life insurance industry working for consulting firms and for some of Canada's largest insurance companies. His particular areas of expertise include mergers and acquisitions, risk and capital management, pricing, planning, forecasting, corporate policies, valuation and taxation modeling. He has worked on many different products, including individual, group and wealth management products. In particular, he was recently involved in developing an ERM program at Canada Life in Canada.

*Copyright © 2005, Society of Actuaries

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

Speaking from a U.S. perspective, it would certainly seem that the Canadian companies, because they're regulated by OSFI and probably also because of the market size, tend to be somewhat ahead of the U.S. companies in implementing ERM. So we're very pleased to have a Canadian colleague with us. Fred is a Fellow of the Canadian Institute of Actuaries. He's also a member of the Risk Management Section Council, the same council I belong to, and he is a member of various other SOA committees and other activities.

Samir Shah is a principal of Towers Perrin out of the Washington office. He works across multiple lines of business and has particular expertise in risk and capital management. Samir is a Fellow of the Society of Actuaries and a financial risk manager certified by the Global Association of Risk Professionals (GARP). Both of these gentlemen have a wealth of expertise relating to ERM implementation.

MR. FRED TAVAN: During an earlier session on integrating and aggregating risks there were two presenters: Mr. Koyluoglu, who has a Ph.D. in operations research; and Mr. Sabatini, who is a consultant in the insurance industry. They both did an excellent job of basically painting a vision of where ERM is heading for the future. But, at the same time, they pointed out the challenges that lie ahead in terms of getting to that future with respect to lack of availability of relevant data, lack of availability of consistent risk measures and the enormous manpower needed to compute economic capital for an entire financial institution.

Mr. Koyluoglu actually shared with us a report from JPMorgan Chase that showed that it now has the capability of computing economic capital for each of the risk categories it has in the risk management framework and then adding those up to come up with the total economic capital for the company. I certainly congratulate JPMorgan Chase on being one of the first financial institutions achieving that milestone.

I'd also like to take a couple of minutes to basically put that milestone or that achievement into perspective. It took JPMorgan Chase 300 man years to come up with that report. If you just assume on average about \$100,000 per man per year, that means that that report cost the company about \$30 million to build. I didn't count all the numbers that were on that report, but, say there were about 20 numbers; each number cost the company about \$1.5 million to build. So there's a lot of investment in getting to that point.

Another way of looking at it is that if your company is ambitious enough to actually decide to try to do the same thing as JPMorgan Chase did and compute economic capital over the next 12 months, that means that you need to go out there and hire 300 people to dedicate to your ERM program over the next 12 months. Now even just hiring 300 people, if you can pull off hiring one person a day, you'd just be busy hiring people for the next year let alone trying to compute economic capital.

Most large insurance companies that I've seen actually have about six to 12 people

allocated to developing their ERM program. So even if we take the high end of 12 people, that means that for 300 man years you're looking at 25 years from today to be able to catch up to JPMorgan Chase. Even if you have a less ambitious goal and say, "We're going to try to get there in about five years." You're still going to need to hire 60 people and allocate them full-time to your ERM program for an average cost of about \$6 million a year. That is very substantial.

Computing economic capital by convoluting multiple distributions with dynamic correlation matrices is definitely a long-term goal. It's not something you're going to be able to do, I think, in the next 12 months or even 18 months. What you need are actually some short-term goals and tools that will allow you to basically start building the ERM framework to start developing an ERM risk management culture within your company and also establishing the communication channels that you need to communicate on an ongoing basis the risk profile of your company up to senior management and the board of directors. Those are the kind of tools that I would like to talk about.

The tools that I'll be talking about are more basic in nature than you would have seen in the techniques that presented in the earlier session. They can normally be implemented within a 12 to 18 month time frame at a much lower cost than something like \$6 million. So I think they're a lot more practical for you to get started on the ERM journey and to eventually get to the goal of being able to compute economic capital.

Now don't get me wrong. I think it's a great goal to have to start on the journey of computing economic capital. It's just I think everybody should realize that if you have limited resources like most financial institutions do in North America with respect to how much is being allocated to your ERM program development, it's not something you'll be able to do overnight.

I'll talk a little bit about the risk management vocabulary, basically some four-letter words you want to have handy in case something goes wrong and your CEO is yelling at you; it would be good to know about these. I'll also talk a little bit about qualitative and quantitative tools that you can use. Again, these are things that you can go out buy such as some software that will help you get off the ground pretty quickly. I'll talk about a generic risk management process that you can use to get rolling.

I'll start off with qualitative numeric rating, and this is something that is usually easy to get going. You can actually start your evolutionary process in terms of moving ahead towards economic capital by starting here and then kind of building on it. Then I'll make a distinction between what quantitative assessment is versus a qualitative numeric rating and get into what I mean by these various words: measure, tolerance, appetite, limits, trigger points and early warning indicators. When you embark on your risk management journey, you'll quickly realize that

some people actually confuse some of these words together. It's very important to clarify what you actually mean by each one across your company.

Here's the first, probably one of the most basic tools that you can use (Tavan Slide 4). On an evolutionary scale, if what we saw in the earlier session in terms of computing economic capital is a monkey, this would just be an amoeba. So it's very, very basic, but it really does help you to start doing some things. To be able to use something like this you actually have to have developed an ERM framework, first of all, so that you can start doing some assessment against that framework. What I mean by framework is you need to have defined consistent risk categories and subcategories across your entire company and make sure that everybody understands exactly what you mean by each risk category and each risk subcategory.

Once you have that defined, then you can actually go to your business unit and your functional experts and ask them to do an assessment on a low, medium or high scale of the inherent risk within each of the risk subcategories within each business unit of each division and then roll it up to your entire company.

The numerics that you're seeing here actually are used quite often, because it's a lot easier to talk to non-actuaries with respect to low, medium or high than to start talking about numbers right away. So this is usually a much easier way to get it started. Usually you're talking about low, medium or high, and after the fact you come back to your desk and say you can't really average numbers and roll things up with low, medium or high, so why don't we just say low is one, medium is two and high is three. Again, this is a very basic way of doing things. Then you can kind of have a scale attached to what you mean by each. The low should be less than 1.49. Moderate should be 1.5 to 2.49, and then high would be more than 2.50.

As you start rolling up, say you want to see within a certain division what the average score is across a certain risk category, or you want to see an average score for an entire division, you can start doing either averages or weighted averages, and you can have different techniques for coming up with your weights.

The next stage in your evolutionary process is to go to something like this (Tavan Slide 5), where you're starting to add some color. Instead of having just a binary system of saying whether your controls are adequate or not adequate, you can now start adding some color. You can have two shades of medium and black and white. There are some things that would fall in the black or white, that's obvious, but some of the things will be gray, and then you can also have some fun in terms of the scale. Instead of having a scale of one to three, you can maybe have a scale that goes from zero to 100. You can try to sit down and attach what kinds of scores should go into each of the boxes. As you see, because low inherent risk and strong controls can produce the lowest type of risk for a company on a net basis, that's what you attach a score of zero to. On the other end of the spectrum, if you have high inherent risk and you have absolutely no controls in place, that's where it will

give you a really high score.

If you try to introduce some basic actuarial concepts into what's going on here, let's start talking about severity and probability and try to bring that into our assessment. You can say if you have a potential formula of both low severity and low probability, now you'll have a desired scale of being low with a numerical rating of one. You can go to a five grade numerical system that just combines the different combinations of low, medium and high on the severity probability. You'll end up with something like this (Tavan Slide 6).

That leads you to something like this (Tavan Slide 7). Now you have severity-probability numerical rating grade where you've attached a score of one to five. So again, the obvious thing is low severity, low probability, and you have a score of one. On the other hand, five and high-low produces a score of three, which is the same thing as the medium-medium and a low-medium. So you can see some consistency here in terms of how you're looking at things. You can now start defining severity as the loss given that the event has occurred. Also, define probability as the likelihood of the event happening in the next 12 months.

Taking it further, you can start providing some guidance to the people who are doing the assessments for you in terms of what you mean by low probability, medium probability and high probability and what you mean by low severity, medium severity and high severity. If you think about it a little bit, you'll see that the scale that you see here is really not that reasonable but just a simple example where we've done a linear scale. In reality what you'll end up with is a much lower number rather than the 33 percent under the low probability end.

What you'll notice when you talk to the non-actuaries in your companies is that they have difficulty in terms of thinking this way. What's easier for them (and this is a tip for you) is to get them to think in terms. For example, if you're talking to an IT person, how long do you think you can go before your system fails? If they say they expect the system to fail once in the next 20 years, well then you can convert that to a probability of five percent. Keep that in mind.

You'll basically develop a scale based on the circumstances of your company. You'll do the same thing with severity. What will be the key driver here is the earnings volatility appetite of your management. Depending on how much of a risk taker or how risk-averse they are, this scale will vary. Again, I just made up numbers for the example.

Now that you've done all that, you can bring it all together and you, all of a sudden, are in a position where you can start building some heat maps. So we brought it together (Tavan Slide 8). The concepts of severity and probability under the inherent risk level with this kind of scale. We also have more shades of gray under the adequacy of controls. When we start looking at this five by five grid, you have a

region in which is low risk, which is indicated by the green area in the region; the really high risk is in the red area; and in the middle there's the medium.

You can have some fun in terms of attaching scores to each one of these grids. As you'll see, it's kind of a weird pattern. When you start thinking about it, you have to make decisions around the threshold points such as where should 18 go, where should 19 go and so on. Different groups of people might come up with different answers, so it would be up to your company. You saw how you can basically take some qualitative assessment approach and then bring it to something that will lead you to heat maps.

The next stage is to actually now start looking at quantitative assessments as you move through, again, your evolutionary process. So for quantitative assessments, we now need to be able to quantify both the amount of the potential loss and the probability of the loss. You're going to define the product of the potential loss and the probability of the loss to be a risk impact, and the unit of that is going to be in dollars. In order to be able to get to that point, you need to identify the various risk measures for each of the risk subcategories, or some people call those key risk indicators.

You're not trying to stay within the principle of consistent risk measures. So that allows you to speed up your process and at least get something based on the information that's already available, as opposed to kind of looking over the world for consistent risk measures right now. You're not necessarily going to need large amounts of historical data. You're just going after things that already exist.

What's a risk measure? The way I've defined it is it's an objective numeric metric that can be used for quantitative risk assessments. Examples of those would be the duration for interest rate risk, quit rate for HR risk and retained face amount for underwriting risk.

Tavan Slide 11 shows examples of some financial risk measures that are available right now; some techniques you can use to quantify some of your financial risk measures. We've been doing a lot of work on the Risk Management Task Force. I'm leading the Risk Metrics Subgroup. I've been building a database of various risk metrics that are currently available out there that you'll be able to use to start getting some quantitative assessments in place in your company.

How do you bring all this stuff together? The table in Tavan Slide 12 is simple, but it actually has a lot of concepts on it. First of all, this would be your ERM framework at the highest level. We have seven different risk categories that are going to represent the various risks that you decide to look at within your framework. Under each one of these, you're going to have various risk subcategories. In total usually you end up, by the time you look at all the risk subcategories under here, with somewhere on the order of 20 to 25 different risk subcategories.

What's been happening as the financial industry has been evolving with respect to risk management is that the banks have focused here, insurance companies have focused here, so all this value-at-risk work that you see on the market and the credit risk has been developed by the banks. Insurance companies have been much better in terms of building models that do simulation. As we've evolved, we're now in a position where there are quantitative techniques that exist for these types of risk categories, which are basically just the financial risks. But we don't have widely known quantitative assessment techniques for operational, legal, regulatory and strategic risks. It's very difficult to find many companies using the same type of approach for measuring operational risk on a quantitative assessment basis.

One of the things that Samir will be talking about later is there is starting to be more and more research in this area. Some of the concepts are being borrowed from the engineering field that seem to have some potential as we move forward.

If you're sitting in a role where you're able to do quantitative assessments and you're not able to do them, how do you bring everything together? You have to go with the lowest common denominator, which means that you have to take some of this information and map it back to the grid that I was showing you earlier so that at least you can look at it on a heat map basis. You would, again, take the severity and the probability information that you would have gathered and go back to the grid that I was showing you earlier and have assessments for your financial risk as well as your operational and other types of risks.

I'm going to talk a little bit now about what I mean by tolerance appetite, limits, trigger points and early warning indicators. Tavan Slide 14 is the conceptual framework. It's something I'm sure most of you can understand very easily conceptually. Basically you set up limits for each one of the risk subcategories. You have the limit at your low point, and you have the limit at your high point. What you would like to do is operate within that tolerance. Now the reason why you define this part is because risk management is not just about always saying no to taking too much risk. It's also about figuring out if you're not taking enough risk. So you need to build this into whatever system you're implementing.

In order to be able to manage within this tolerance level with some level of confidence, you need to also define some trigger points at which you're going to start taking some action. We have, therefore, limits. We have tolerance levels and we have trigger points. I'm just going to give you a little bit of definition around each one of these. Tolerance levels are acceptable levels of risk for a specific risk category as defined by the company management in terms of loss amounts, error rates or other types of units. Again, you're not trying to stick to a consistent risk measure because it's too difficult at the beginning. What you would like to do right up front to get the journey started is to grab whatever information is available. That is a lot of work in itself.

Risk appetite is really interesting, because I bet you that if a lot of you went to your

board of directors or your senior management, and asked them to define what risk appetite is, you'd have a lot of difficulty getting an answer out of them. By definition, basically what you mean here is the level of aggregate risk at which the company can successfully manage its business for an extended period of time.

Risk limits, as I said, are the boundaries of tolerance levels above which too much risk is being taken. Again, that was just the high end of it, but you can also talk about the low end.

The key concept with trigger points is the trigger points have to be predefined so that everybody knows what they are up front before you start doing anything. They're known up front. Once they're activated, you have predefined steps to be followed by predefined people. It's almost like when the fire alarm goes on the firefighters know exactly what routine they will have to go through to get to the site. There are no ifs or buts about it. They don't have to sit back and have a discussion before they jump on the wagon as to how to do it or what to do when they get there. These trigger points can be based either on risk measures or on early warning indicators. You're just trying to make it as simple for yourself as possible in terms of using information you already have available.

Early warning indicators are readily available objective metrics that provide information on direction and/or magnitude of a specific risk. They're not necessarily risk measures as we previously defined them because they're not going to be necessarily in terms of severity or probability. An example of that is something like using unrealized gain and losses from market risk. It's just an early warning indicator because most companies obviously have to calculate that just for accounting purposes so that it should be something that's available to you.

What's the process now that you can use to bring all of these things together? Tavan Slide 22 shows a generic process. For each of the risk subcategories that you have within your framework, you pick a certain measure that everybody agrees across your organization is going to be the measure for that specific risk subcategory. You determine the appetite for that risk measure. An example of that that most people are used to already is a lot of investment policies have a duration of, say, a half a year as the limit and their mismatch between their assets and liabilities. So that's kind of what I mean by that process, but applying it not just to interest rate risk. Now you're applying the same concept to all the other risk subcategories across your framework.

Once you determine your appetite, you set tolerances around that. Once you have your tolerances, you set limits and define your trigger points. You define actions to be taken when the triggers activate and then you assign ownership. So now you have something that's predefined and something that could be applied on a consistent basis. You can certainly start feeling a little bit more comfortable that your risks are being managed not on an ad hoc basis, but on a systematic basis.

How does that generic process work in practice (Tavan Slides 23 and 24)? You basically start off by selecting the risk measure, as I said. Then your senior management should take a look at the capacity curve of that particular risk and their willingness to bear the risk to determine the risk appetite for the company. Once you have a feel for what that appetite is, you can now start setting some high-level risk tolerances around that appetite. Once you have the tolerances, you can set some limits. Once you have the limits, you can start setting some of the high-level trigger points.

Now in some cases the high-level risk tolerance will be usable at the lower levels in the company. What you start off with is something at the corporate level, but then you need assessment and tools that can actually be used at the divisional level, so some of those things will be directly transferrable, some will not. If they're not, then you have to basically convert the lower level risk tolerances using the high level as a guide, but converting to something that the divisions can use.

Once you've gone through that, you can now choose low-level tolerances that the business units will be using. You use that to set low-level risk limits. Then you define your low-level trigger points and you define actions to be taken when the trigger is activated and so on.

Here's a quick example of how that would work (Tavan Slide 25). So if you look at your reinsurance processes, for example (and this is not necessarily all of the risks that your insurance operation is taking), say you're just looking at the counterparty credit risk and the operational risk of unsigned treaties. Again, I just want to show that you can take any one of these processes in your company and go through this generic kind of process and break it down so that you can put some systematic processes in place.

So for credit risk you start out saying the risk measure we're going to select is a rating agency credit rating. Senior management decides that we're going to provide some high-level guidance in terms of the credit risk limits, and we're going to define those in the reinsurance policy that we have. The tolerance that we're going to have is not doing business with anybody or any counterparty that has a credit rating less than A-. You keep going down and you get to a point where you decide, well, okay, my trigger point is going to be the same thing as my risk limit. They don't necessarily have to be different, so in this case I've just kept it simple, but it could have been something that would have been higher even than A-, so that you used that as a trigger point.

This is one of those examples where it actually is transferrable to a low level, so your high-level and your low-level tolerance and trigger points can actually be defined on the same basis. You can basically go to your division and say, well, I don't want you to be doing any business with a counterparty that has a credit rating of A- or less. You just basically bring that down, and you assign it to the actuarial vice presidents in your company. It could be a different person depending on your

circumstances. The action that they have to take is to escalate that to the risk management committee for discussion so you can figure out what to do. Again, it's a very systematic approach. Everybody knows what to do. It's not on an ad hoc basis.

There's a second example (Tavan Slide 29) that I provided just to give you a feel for a case where it doesn't easily transfer down to the low level, and you have to keep changing as opposed to keeping the A- like in the previous example. So for operational risk of unsigned treaties what you decide to do is say, if I look at just the net amount at risk (NAAR) on my unsigned treaties, I'm starting to get kind of uncomfortable if I have more than \$100 million of treaties that haven't been signed yet. So you say my risk limit is going to be \$100 million of unsigned treaties at the same time outstanding. You set your trigger point to be \$85 million, just to give yourself some room to take action before you actually get to \$100 million.

You can't just use the \$85 million because you probably have several divisions. You can't give that \$85 million limit to all of your divisions because, obviously, if they each have \$85 million the likelihood is that at the corporate level you're higher than \$85 million. Now you have to convert that into something less than \$85 million that you can allocate out to your division. You convert it to a lower level tolerance, say \$50 million in unsigned treaties per division.

The trigger point now, again, is set even lower at \$45 million so that the division is going to actually take some action before they get to its limit. The action to be taken is immediate signing of treaties. It's never that easy. You can't just go out and sign them, but at least know that it's showing up on some kind of radar screen somewhere, and you know that it has become a high priority all of a sudden to make it at least a special project and to make sure that you are getting them signed and you're not getting exposed to too much NAAR in unsigned treaties. Then you assign that task to your divisional CFO to make sure that somebody is actually taking care of it.

MR. MUELLER: As Fred was saying, if you go to the Society's Web site and go into "Sections," there is a page for the Risk Management Section. Under "Risk Management", you find "Risk Management Task Force". There are different topics, and one of them is actually "Risk Metrics". The Risk Metrics website, which is just one page, does have some of the tools listed as an Adobe.pdf file to download and print out in much more detail than what we have in the presentation.

Fred has given us a lot of tools and analytics to set up an ERM process. Samir is going to talk to us about some of the newer and non-traditional methods of how to define and quantify operational risk.

MR. SAMIR SHAH: As Hubert mentioned, I'm here to continue our discussion on analytics by focusing on operational risk. I'm referring to operational risk in a very generic sense. I'm not referring to a specific definition like the Basel definition. It's

just really all the messy risks that are not market risk, not credit risk, not insurance risk, not liquidity risk.

I will be focusing my comments on several modeling methods that have been borrowed from the engineering sciences that I think offer a lot of potential in quantifying operational risk. I'm also going to briefly talk about coherent risk measures. I think it's important as it relates to our discussion of the modeling of risk.

(Shah Slide 1) The first thing is before we actually talk about the modeling approaches, let's first understand the fundamental characteristics of operational risk so that we can figure out what kind of modeling approaches we should apply. There are a couple of things that stand out for me. The first is, of course, operational risks are very much specific to the company because they arise from the processes, the technologies and the organization of the company. What that means to me is that you can look at industry data, but it's not going to be as reliable as it relates to assessing your own risk. You can look at your own data, but you have to also be careful that the data that you're using is representative of your current operational environment. We know that operational environments change all the time. There's always something going on in a company. There are changes in technology. There's Six Sigma and new business initiatives. That's one thing to keep in mind is that our data has to be representative of the current and future environment.

The other thing is that we're hopefully not measuring operational risks just for the sake of measuring it. We're focusing on that to figure out how to manage it. What's important is that we need modeling approaches that will allow management to look at risk and figure out if we make some changes to operations to hopefully mitigate that risk. We have some idea of what the impact will be. What's really needed is some sort of a decision tool for management to decide between alternative options in managing operational risk or, in fact, managing all risks.

What that says to me in terms of the modeling approach is that we need an approach that's a causal approach rather than a statistical approach, because we need to be able to say that if you do this, this is how the risk changes and here's the cost of doing that. Then you can compare the effect.

I said I was going to talk about was coherent risk measures. The reason for that is because operational risks are typically not normally distributed. They usually have very fat tails. When you're combining risks for allocating capital against these risks, you have to make sure you're using measures that make sense.

(Shah Slide 2) I think there are actually many different modeling approaches, and there are lots of different ways to organize these approaches. What I find useful is to organize them along the dimension of the extent to which they rely on historical data and the extent to which they rely on expert input. I think for operational risk, because we're going to eventually have some data gaps, we're going to have to rely

on expert input. As you saw in Fred's presentation, a lot of those numbers obviously come from sitting down with experts and getting their feedback on whether risks are low, medium or high.

The methods on the left side are the ones that, as actuaries, we know very well because we apply them in modeling of traditional risk. Those are all statistical in nature. The ones on the very right side are used primarily in the decision sciences, and they rely primarily on expert input, almost totally on expert input. But I think the methods that are the most interesting are the ones in the middle where you have some data. You may have more data in the future, but you really need to supplement that with expert input. I'm going to focus on system dynamics simulation, Bayesian belief networks and fuzzy logic. I'm going to do that through example. I won't get into the technical merits of the approaches. This is intended to be kind of an instructional session, but I can refer you to places where you can get more technical detail.

(Shah Slide 3) First is systems dynamic simulation. This method was developed maybe 40 or 50 years ago by Jay Forrester at MIT. It was developed in engineering sciences. What's amazing about this approach is that it's slowly migrated into other applications. In the last five or 10 years in particular, it's become very prominent in addressing general management issues. It's taught in all the business schools. Many of the business schools have executive sessions on systems dynamics modeling and systems thinking. The reason it has spread is because it's a very robust approach for modeling any kind of complex situation. So I think that it's worth looking into as it relates to operational risk.

The other area, which is somewhat similar, is Bayesian belief networks. It's not quite simulation. I think of it as kind of a discrete alternative to systems dynamic simulation. But it's an approach that's being used more and more often in the banking sector so, again, I think it offers some potential.

The last approach that I'm most interested in is fuzzy logic. I'll relate that to the kinds of examples that Fred had when you sit down with experts and ask them to rate things on low, medium or high.

First, let's talk about systems dynamic simulation. You would go through kind of system mapping (Shah Slide 4). There are two steps. The first step is to draw a system map. A system map is simply drawing on a piece of paper the cause/effect relationships of all the variables in the system that you're modeling. Let's focus on the event of IS failure, which is depicted in the middle of that map. There's a little circle that says IS failure. The way to apply system mapping is to first go back and say, what are all the things that could cause IS failure? Then go forward and say, what are all the things that it could result if you had IS failure?

As an example, if you went backwards, you'd say IS failure could be caused by some kind of virus infection or some kind of a Web site hack. A virus infection, in

turn, would be caused by employees not following policies or maybe you're not updating your virus protection software. The extent to which employees follow policies is affected by how many resources you commit to that area. The same thing is true with the Web site hacking. If you go backwards and say, what are the key factors that would create the likelihood of a Web site hack? You realize it's a technology. You need to look at what kind of firewalls you have. How popular is your company name? What kind of brand recognition you have? To what extent do you use your Web site for services?

Now if we move forward from IS failure and say that the event happened, what are the consequences? The consequences may be the e-mail system is shut down. Desktops and servers are shut down. The impact of that is that it takes time to recover those systems; you end up losing some information. How much time it takes to recover and how much information you lose depends on things like to what extent you're appropriately staffed, what kind of processes you have to backup things, what kinds of relationships you have with clients and how much you could lose in terms of penalties and so forth. The point of this is that it's one thing to just assess IS failure as high, low or medium.

We'll see that some of these circles actually have a little bar in them. They refer to actual operational decisions whereas other circles are actually the consequences where you could actually get a distribution out of this. So the first step is mapping this.

The second step is to quantify this map. By breaking it down into little pieces, it's much easier to quantify it. Some of these things you can actually use data for and some of these things you may have to use expert input for. So, for example, the impact of brand recognition on Web site hacks may be something that you can get industry data on. But the impact that employees following policies has on virus infection may be something that you rely on expert input on. There may be some things where you may have data. There may be some things where you have to rely on expert input and you kind of bring it together. Even the areas where you rely on expert input, it at least tells you if you don't want to continue relying on expert input because this is a critical risk. It focuses your data-gathering efforts as well. This is systems dynamic simulation.

The other method I mentioned was Bayesian belief networks, which is very much like a system map. I've taken part of that prior system map and blown it out (Shah Slide 5). For example, virus infection is affected by frequency of virus protection updates and employees following e-policies. The essence of Bayesian belief networks is to, based on either data or expert input, develop conditional probabilities. The frequency of virus protection update (the blue box) is actually a management decision so you can decide whether you do that every day, every five days, or every 10 days. You can put a "1" corresponding to your decision. The extent to which employees are following e-policies may be based on data that you can gather and then you have to sit down with an expert or, perhaps, industry

data to develop conditional probabilities on the likelihood of a virus infection or the severity of a virus infection conditioned on the extent to which employees are following policies and the extent to which you're updating your virus protection software. I think of Bayesian belief networks as kind of a discrete approach to systems dynamic simulation. Again, this is something that's being used more and more often in the banking sector.

The next approach is the fuzzy logic modeling. I'm going to walk through the steps of actually trying to use fuzzy logic modeling to assess market conduct risk. The first step is to identify key risk indicators (Shah Slide 7). As an example, let's say that the losses due to market conduct risk are affected by the experience of agents, the product complexity and annual growth rate. Years of experience in product complexity may be obvious. With the annual growth rate, depending on how fast you're growing, there may be pressure to sell, which could create some risks.

The second step is to create fuzzy variables. This is where Fred had talked about specifying low, medium and high. The essence of fuzzy logic modeling is to use linguistic variables. The difference between linguistic variable and numerical variable is instead of assigning values to it, you assign words to it. The variable is agent years of experience, and the values it can have are low, medium and high (Shah Slide 8).

You can sit down with an expert and say from two to five years is low. From five to 15 years is medium and 15 years and above is high. But the way that we think is not so much in these discrete areas by saying, well, 4.5 is low and 5.5 is medium. For us, these things overlap so what fuzzy logic does is it creates a system which facilitates getting input from experts and by allowing you to gather information the way we think rather than how computers think. In this example, somebody who has five years of experience is to some degree somebody with low experience, and what's on the Y axis is the degree of membership in that value. For example, in the yellow line there the five years of experience would say that this person is low with a degree of 0.3 and is medium with a degree of 0.7. You have to calibrate these variables, again, similar to in Fred's approach. You have to calibrate these variables to the actual values of agent years of experience in a way that's reasonable and an approach that people would agree to. Now when you do that, then you can ask them to provide some rules that relate the key risk indicator to the losses that you're trying to model. Basically, they help you quantify this relationship between the causes and effects.

(Shah Slide 9) So, for example, if you sit down with experts you can see them saying things like if years of experience are low and product complexity is high, then risk is very high. Now if years of experience are high and growth rate is poor, then risk is medium. This is how people think. This is how people naturally can provide input without having to contort their thinking into the way we quantify these things. Basically by talking to you, they can specify these rules.

In fact, I know fuzzy logic has been used in insurance underwriting. I know it's usually used on the P&C side. I don't know to what extent it's used on the life side, but I don't think it's foreign to insurance.

(Shah Slide 10) Once you specify these rules, then math takes over. The first thing to do then is to take the key risk indicators and develop distributions for those. For years of experience, you can take a look at all the agents and you can develop a distribution based on the existing agents. Product complexity is something you can rate on a scale of one to 10 and map all the products against that. Expected growth is something that you can apply a distribution to. If you can apply a distribution to the key risk indicators and you have those fuzzy rules, there is a fuzzy calculator that will take those things and develop a probability distribution for loss. I won't get into the math, but there is essentially software that you can get to do that. The output of this is actually a loss distribution, and it's actually a multidimensional loss distribution.

(Shah Slide 11) This is actually a screen shot from MATLAB. MATLAB is a tool that's used often in engineering sciences for modeling. I am using the fuzzy logic module to develop a toy model. This shows on the two axes on the bottom. One is product complexity from zero to 10. The other one is years of experience going from zero to 15. On the vertical axis it's the market conduct risk measured in terms of expected loss going from zero to 60. So you can see what comes out of a fuzzy logic calculator is a risk surface and you can see from that the impact of years of experience on risk is relatively low when the complexity is low. As the complexity increases, the impact of years of experience is much greater. This is something you can go back and sit down with experts and say: Does this seem reasonable?

I think this process lends itself to the kinds of qualitative approaches that Fred was talking about. When you sit down with somebody and ask them to say low, medium and high, you can get the same kind of input and just convert those variables into fuzzy variables. Then you can also ask them for these kinds of rules that relate that to the loss. You can actually develop some distributions on that. That's on modeling approaches.

(Shah Slide 13) The next thing you want to talk about is coherent measures, and the reason for doing that is because, as I mentioned before, operational risks will tend to have very skewed distributions. They tend to have very fat tails. The traditional measures that actuaries use may not be appropriate if what you're trying to do is assess the risk and then apply capital against those risks and to allocate capital or attribute capital to those risks. I'd like to go through an example that shows that some of the measures like value at risk (VAR) or probability of ruin or even things like standard deviation are not coherent when you use them for capital allocation.

Let me just go through this example (Shah Slide 14). There are two different companies. There are three different scenarios, whatever the risk is, and the

scenarios are very similar. Scenario one has 97 percent probability. Scenario two has 2 percent probability. Scenario three has 1 percent probability. The main difference is the outcome associated with scenario three. In Company A, scenario three generates \$28,000 in loss. Scenario three in Company B generates \$55,000 in loss. In one and two, the losses for Company A are slightly lower. Now the expected loss, therefore, is the same between these two companies.

If we were to use a VAR-type approach and say that we wanted to set capital at 99 percent, then for Company A we would put aside \$10,000 to cover the 99 percent likelihood. For Company B, we, again, would put aside \$10,000. Let's ignore the economic cost of ruin (ECOR) and ECOR ratio column just for a second. I'll come back to that.

First of all, Company B is much riskier, but given our threshold, given our objective, they appear to be the same. What happens if you combine the two companies? (Shah Slide 15) If you combine the two companies and you just convolute all the combinations, you can see all the joint probabilities. If you cut off at the 99th percentile, what you see is that what you actually need. If you look at the loss column, the \$38,000 is the cutoff. Everything below that adds up to a 1 percent probability. So, if you wanted to use, again, 99 percent VAR, you would say you need \$38,000 to combine the two to cover the combination of Companies A and B. This really doesn't make sense because for Company A we only needed \$10,000. For Company B, we only needed \$10,000. We know that when we add up companies that are not perfectly correlated, you should need the same or less. You should need \$20,000 when, in fact, you need \$38,000. This is a stylized example just to emphasize the point. This kind of thing happens, particularly when you have fat tails.

(Shah Slide 16) Whether using VAR or any kind of coherent measure, you see this kind of an effect. There's a very important paper that's noted at the bottom of Shah Slide 17 that set up some axioms for what would be a coherent risk measure. It's a very technical paper, but the essence of it is that you have a coherent measure if it follows these four criteria. (Shah Slide 17) The first one is that if you have two portfolios and one portfolio is always better under all scenarios, then you should need less capital for that portfolio, which makes sense. The second one is that if you're combining uncorrelated risk, as we were in this case, then you should never need more capital. If you're combining perfectly correlated risk, then you should need exactly the same capital requirement as the sum of these two. The final criterion is that if you add some non-risky elements to a portfolio, then your capital should be reduced by that much. These seem like very reasonable criteria, and it's amazing that things like VAR and, as I mentioned, standard-deviation-type measures don't meet these criteria. Probability of ruin and VAR actually are not coherent because they fail the second criterion, which we just showed in the example. I think that standard deviation, for example, fails the first criterion.

There have been different types of measures that have been suggested that are

coherent. One of them is what we call ECOR ratio. There are other tail VARs that have been defined that prove to be coherent. Let me go through the ECOR ratio, but I don't want to suggest that that's the only one. (Shah Slide 18) The ECOR is the expected loss given ruin, and that's divided by the expected value of that distribution to get a ratio. First of all, under the VAR approach, we saw that Company A and B needed \$10,000 each. The combination needed \$38,000 instead of \$20,000, so, in fact, there was a diversification penalty. But under the ECOR ratio, Companies A and B would need 15 and 42. When you combine them, you actually need less money, which is what you would expect.

As I've mentioned a couple of times, when we are specifying different risk measures, and Fred, you were saying that at this phase we're using different risk measures for different things, it might be difficult to come up with one risk measure that works with everything. I think it's okay to use those other measures. If you're using it for capital allocation, then I think you need to be careful.

MR. MUELLER: How many of you actually have an ERM system or process at your company, or as a consultant have worked in that field? (About half of the audience member raised their hands.) Some were cautious, some weren't.

What are your experiences with ERM? What are the problems you're facing? What are some issues that were on your mind when you came to the session that you're trying to get answered?

MR. MATTHEW S. EASLEY: Fred, on some of the examples of the semi-numeric techniques that you were giving for operational risk, have you ever seen anybody try to use those in more traditional risks where the challenges of modeling have proven to be a problem, thinking about the 300-man-year process to get something done?

Certainly one of the things we're facing is the challenge to get all of our systems doing things. Even if we know how to do it, it takes a lot of time and money to get there.

MR. TAVAN: I think, again, you're forced to go back to the lowest common denominator when you're trying to bring everything back to common platforms. So even though you might have quantitative techniques that help you quantify in dollar terms your financial risk, if you want to bring it back to common platform you need to bring those risks back to that matrix I was showing you where you bring it back to severity and probability of those financial risks.

You should have certain measurement techniques that you already have in place and be able to use those to get a sense of what your severity and probability are around those types of financial risks. Take that and map it to the grid and be able to at least attach some kind of risk score. What you end up doing is you actually, instead of coming up with economic capital, you're building yourself a risk index for

the company, and that risk index within the heat map actually allows you to see what's happening to the risk profile of your company.

At least you can do some directional analysis from year-to-year in terms of whether your risk in that particular category is going up or down. What's happening to certain divisions? Are they taking more risks, fewer risks? At least in the 12 to 18 month time frame, you can start doing that kind of an analysis whereas, as I was mentioning, to get to the stage where you actually have economic capital calculations that lead you to the next phase of analysis will take you much longer.

MR. MUELLER: One of the things I've seen, just taking what Fred talked about to the next level, is, if your company has, let's say, 20 different business units and each business unit is looking at identifying the top 10 or 20 risks by severity and by frequency, you can then also aggregate those at the company level by then combining those 200 or 100 into the top 10 or 20 risks at the company level. In particular, some of the Canadian companies have done that.

MR. TAVAN: Yes. Another thing you can use for that kind of information, again, is to prioritize the top 10 risks and it gives you a systematic approach that you're applying consistently from year-to-year. So, it's much better to have that than to have nothing at all while you're waiting to just get to the economic capital point. It makes a lot of sense to have that in the interim.

MR. ROBERT J. TIESSEN: Fred, one of the things that we've come across in trying to measure the probability of error in one of the examples you had is if you have a process underwriting your claims where you do thousands of events per year and you're trying to determine what the likelihood of an error is. Say you have 100,000 cases, and so, if you make 10 errors, does that mean the probability of error is 100 percent because you made some errors? Or is your probability of error 0.01 percent or whatever it works out to be? Is that something that's an issue in your process as to how you measure the error rate with these huge processes that have thousands or hundreds of thousands of examples?

MR. TAVAN: I think what you need to do is figure out if you have 100,000 cases and you end up with 10 errors, did those errors actually lead to some kind of loss? I think what you want to distill it down to is to only focus on those ones that actually lead to some kind of dollar level loss. The example you described is actually one of those things you probably have some kind of data on. If there was an error, how much of a loss was associated with it? That will give you a sense of how to approach the probability angle as well.

MR. SHAH: It seemed like the risk that you're describing, because you have a lot of data with it, can be modeled in a strictly actuarial sense by doing frequency severity type models rather than relying on expert input.

MR. SANDEEP KUMAR: Fred, you mentioned a particular initiative would take

approximately 18 months. Could you give a little bit of feel for what those 18 months would be comprised of in terms of how this goes? Could you give a sense of what kind of expense would be involved?

MR. TAVAN: These are just my opinions and not what happened at Canada Life. As I mentioned in my presentation, the first thing you need to do is to really build your ERM framework. There are several out there you can borrow from right off the bat, so it depends what you want to do. To really custom make one will take you a little bit more time. You could just borrow one of the ones that are already available.

The one that I had on the presentation is the framework that's used by the regulator in Canada—the Office of the Superintendent of Financial Institutions (OSFI). I think it's widely available. You can borrow it. There's one framework that the American Academy of Actuaries has developed so you can just take that and start off running with it. If you just take one of those, it will save you a lot of time. What you'll have to focus your time on is actually communicating that out to the rest of your organization and making sure that everybody agrees with those definitions. You will, undoubtedly, run into some complications with respect to some people not agreeing with some definitions. It will take you some time to work through that. I wouldn't be surprised if you spend a good six to eight weeks even if you borrow something that's already available to get that running.

In terms of resources, it depends on how big your company is and how many people you have. You usually get started in the corporate area with it, if you're just starting from ground zero probably with one to two people at most in terms of resources. So say if you have two people working for a whole year on it to get it off the ground, you're looking, again, at an average cost of \$100,000 per head, about \$200,000. You will have to spend some time developing some software or buying some software. There are some available now.

Once you have your risk management framework defined and communicated out, the next stage is to find the experts in the company that will actually help you make the assessments. That will take you a couple of months as well in terms of identifying all those people. The bigger your organization, the more time that's going to take. Then you're going to sit down with them and go through the assessments of each risk subcategory. Normally that will require some kind of database where you're going to need to track all that on a consistent basis and have some kind of methodology such as the ones I've shown to kind of bring it back together and roll it up, consolidate it either by line of business or by risk category so you can slice and dice your risk profile in different ways.

That journey in itself actually teaches you a lot about what kind of risk appetite your company has, how different people feel about risk, building a common language across the company. So even though you're not doing some of the fancy mathematical things, you're actually putting a lot of good things in place with

respect to developing your own risk management culture.

Again, you don't have to make it very complicated. You can stick to the first kind of assessment tool I showed as your first time around. Then the second time around, it just moves to the more sophisticated lines. Once you reach that level, the next stage is mapping out your evolutionary step, I guess, in terms of how you want to get from there. Each year improve on that because you don't want to be doing qualitative assessments forever. You want to move more and more towards the quantitative.

By mapping it out, at least you're breaking down the problem into steps that you can actually take, because at the beginning it's just overwhelming. There are so many different risks to look at. If you just try to have a single measure for all of them, you're going to go crazy pretty quickly. So you need to break down the problem. Once you've broken it down, at least into different risk categories, you can have a parallel approach where you're tackling each one of those risk categories with different teams of people. So, again, if you have an ambitious program of getting to economic capital quickly, you can start building your teams that way and develop each one of those quantification techniques and parallel at the same time as having a vision that will eventually bring you back together.

MR. MUELLER: One thing to add to this discussion from our expertise in helping companies implement ERM is that depending on the complexity of your organization, it may be a challenging and daunting task to try to have an overarching framework that works equally well for everything. What might work is to set something up with buy-in from senior management. Communicate to them what you want to do, develop the risk appetite and then start to look for some problem areas in the company and try to fix the key problem areas first. Try to get some quick wins that will give you more credibility when you roll out this process into other groups.

The ones I've seen fail are when a company is very methodical about it, maybe has this overarching framework, and tries to do everything parallel in every single entity of the company. Those things usually fail because they are just too cumbersome and too complicated.

MR. SHAH: I was going to add that it seems like there are instances when companies have kind of a top-down approach to ERM. There's been established support for ERM from the top, and that's fine to kind of go this route. There are many instances we've seen where there are few people that really support it and a lot of management has not bought into it. They think oftentimes that this is just doing what we already do better. I mean it's just good business. It's not something new. Also, they see it as a cost. How do we actually save money? In that case, I think you need an approach (which is what I think Hubert was getting at) that finds the one or two problem areas, focus on that and see how that could be extended to broader initiatives.

In terms of failure, we've seen where companies have actually created the chief risk officer position, and then a year later it was gone.

MR. MUELLER: It will also help you define risk appetite for the company, because you can talk to senior management or to the department heads and ask: What are the things that you'd like to avoid? What are the risks that you can't afford to take? What's the financial hit that you can't take? Is it the ratings downgrade? Is it a loss of more than X million dollars in a given quarter or line? That will quickly get to the bottom line of what risk appetite is for the company and for the business unit.

MR. TIESSEN: I have a question for Samir about fuzzy logic. Have you actually worked with companies that have put that in place? It seems to me that it could be something that could never end. You could keep coming up with rules. You gave the example about driving and you could ask if the person had a driver education test, etc.? Is there some need to put boundaries on that, or how does that work?

MR. SHAH: Yes, absolutely. I think, in general, you need to put boundaries on which risk you quantify, first of all. The top 10 approach, I think, lends itself to this kind of thing. If you take the top 10 risks and specify their key risk indicators, first of all, it's important to do that to see if you have common risk indicators across many of these risks. It also gives you a method for seeing if there's aggregation of risk across. But I think if you focus on just those 10, this is a doable thing because, first of all, you're already going to experts and gathering this information. You're sitting down and asking them, in many of these cases, is it low, medium or high?

The only thing in addition you need to do is to, first, calibrate those low, medium and high. Again, this is something that I think you do. You say 0 to 33 percent is low. So there's not much more work beyond what's being done. I think it's just a better way to use the information that's already being gathered if you can put it in a fuzzy logic framework, because then you can ask them to specify these rules and then you can crank out these results.

So to answer your first question of whether anybody is actually doing it, I don't know if companies are doing it. We're talking to a company right now in applying this approach for exactly the top 10 risks kind of framework. I keep seeing a lot of potential for it, but I think maybe it's just too weird and too foreign for companies to be interested in it, but I think it offers a lot of potential.

MR. MUELLER: The interesting thing is, as actuaries, we tend to model financial risks because we can model them well. Stochastic modeling capabilities have increased, so everybody's getting pretty good at stochastic modeling. But if you look at the failures, most of them have been because of operational risks—rogue traders, market conduct—all those things bring companies down. It tends to be non-financial risks that really bring companies down. That's something to keep looking for. If you look at the Basel II framework that the banks have put in place,

they're much more focused on operational risks as well.

What's interesting to me on the Basel framework is that even the banks only quantify some pretty vague levels of capital that operational risks should be as a percentage of the total capital, something like 12 to 15 percent of capital. In your experience, have you ever looked at when you've quantified operational risks and you developed the capital to hold for that operational risk, how much was that of the total company capital that you were holding or were supposed to hold?

MR. TAVAN: I'm still at the research and educational stage on operational risk, so I can't comment on that yet.

MR. SHAH: I don't know if we quantify the capital for that separately because where we've quantified operational risk it's been in aggregation with the financial risk. It's been quantifying it in this way and putting it into kind of a stochastic model to look at that aggregate impact. But I think it's very much specific to companies. The areas which gave the greatest source of risk in areas where we've worked have been in distribution. I think that when companies have been modeling product risk, reflecting the risk associated with the distribution channel is actually something that was surprising.

MR. MUELLER: Yes, it can bring large companies down, as we've seen most recently. Marsh is suffering from that right now.

MR. KENNETH A. LASORELLA: Do you have confidence in any analytic tool that does a good job, something very practical that could aggregate risk? Specifically if we have two risks, A and B, and they might be seemingly virtually uncorrelated in a general sense, but once we move toward the tail if things go bad for A, let's suppose things also go bad for B, or at least a high probability. I know there's copula that's out there. Do you have any comments on a practical tool?

MR. SHAH: First of all, I should mention that on the Casualty Actuarial Society (CAS) side, there's a task force that is going to be presenting some results on correlation independencies, and I think it's going to happen some time in the spring. So you may want to go to the CAS Web site to look at the task force. But those use copulas and are very technical approaches.

I think that the simpler approach, again, is using these causal approaches. If you map the cause/effect relationship, just sit down with experts and draw the arrows, I think you will find that there are areas which affect multiple risks. Without even modeling it, I think it gives you some indication of where there's concentration of risk. In all the approaches that I talk about, any causal approach will help you do that. Whereas I think the approaches that the task force is focused on are statistical approaches. You will get a measure of aggregation, but you won't know why.

MR. MUELLER: It's also worth mentioning that the CAS and the SOA are working

closely together in the area of risk management and ERM because, obviously, a lot of those risks go across sectors. There are quite a few multiline companies in the market.

DR. HARRY H. PANJER: I just want to make a couple of comments with respect to modeling aggregation risks. Really what you've done is outline two approaches; one is a causality, what I call a structural approach, where you essentially build a structure of the relationships into the model and then the correlations come out automatically.

But the copula approaches are extremely useful when you have an underlying belief that if something goes bad, everything goes wrong. You don't really know why. You don't really have an underlying driver, but you have this underlying belief that things can go wrong. When some things can go wrong, other things can really go wrong as well without describing it. The copula approaches are actually not that technical and not that difficult to apply. The group that has been working at the International Actuarial Association (IAA) on solvency (Stuart Wason was chair and I was a member and a number of other people from North America were members) actually illustrates some results using a T copula. It is actually very easy to apply, and you can actually calibrate the degree of dependence. What that allows you to do is actually test the sensitivity of the result to the dependency, so dependency can then become a variable that you can use for sensitivity testing. Even if you don't know what the degree of sensitivity is, you have some hunch that there may be sensitivity. It provides a very useful tool for sensitivity testing of dependency as a variable.

MR. MUELLER: Right. In fact, just to point out, again, if you go to the SOA Web site under "Risk Management", there's a section on Economic Capital, and there's an Economic Capital Specialty Guide. In the Economic Capital Specialty Guide, you'll find a very good literature section that has a discussion of papers on this topic, including the one from Artzner, Delbaen, Eber and Heath that was mentioned earlier and including the one on copulas that Harry just mentioned. In addition, a similar Specialty Guide on ERM is currently being drafted. There's a wealth of information out there, but you do have to do a bit of digging to get there.

MR. GEORGE E. SILOS: My question has to do with the Rolls Royce, JPMorgan, Citibank approach. Is that decision really a binary decision that you either have to decide to spend 300 man-years on that and get that report with a high degree of confidence? Could you try to put together a similar report, have less confidence, but do it in three man-days, or something in between? Is too dangerous to try to do it a little more cheaply and quickly?

MR. TAVAN: My opinion would be that it's too dangerous because there are just too many different types of risk. If you're trying to do it in three man-days, you're going to overlook a lot of things and the number that you're coming up with is just oversimplifying what's going on.

MR. SHAH: Yes, I think there's danger in oversimplifying, but I do believe that the approaches that work have been approaches that you phase in this stuff. I don't think you were implying that you have to do it all in kind of a big bang. In fact, that doesn't work. I think you're right. You can do it, but the idea is you have a framework and to have the end goal in mind and then the pieces are going to fit into place. I think that it's okay to do that.