RECORD, Volume 29, No. 3*

Orlando Annual Meeting
October 26–29, 2003

Session 126PD
Global Developments in Insurer Solvency Assessment

Track: General Session, Financial Reporting, Investment

Moderator: STUART F. WASON
Panelists: ELIZABETH S. BRANAUM
HARRY H. PANJER
DAVID K. SANDBERG

Summary: Attendees at this session gain an understanding of risk, not just for solvency assessment but also for risk management and performance reporting. The panelists explore the proper design of insurer risk models.

Attendees learn about:
- *International actuarial initiatives to help develop a global framework for insurer solvency assessment*
- *Key insurer risks*
- *Key elements of risk (volatility, uncertainty and extreme events)*
- *Techniques for modeling insurer risks and their dependencies*
- *The selection of an appropriate time horizon and confidence level*
- *A case study of both factor-based and internal approaches.*

MR. STUART WASON: First up on the agenda is Liz Branaum, who is a technical actuary at USAA Life Insurance Company in San Antonio, Texas. Her current responsibilities involve enterprise risk management and company experience studies. Liz has over 20 years of experience as an actuary, with a primary focus on pricing and product development of various life insurance and annuity products. Liz is a fellow of the Society of Actuaries, member of the Academy and certified life underwriter. She received her B.A. in math from Trinity University in San Antonio, Texas.

Note: The chart(s) referred to in the text can be found at the end of the manuscript.

David Sandberg is second up on the agenda today. Dave is second vice president and corporate actuary at Allianz Life of North America. Dave is very active with a whole variety of Academy solvency and financial reporting committees. He is also a member of the International Actuarial Association (IAA) Insurers Solvency Assessment working party that you're going to hear more about in just a few minutes.

Harry Panjer is our third speaker. He is a professor at the University of Waterloo, the immediate past president of the Society of Actuaries and a past president of the CIA. Harry is also a member of the IAA working party.

MS. ELIZABETH BRANAUM: Basically I'm going to talk about insolvencies in general. My examples are from the United States, but the things that happened in a lot of these companies could be very international. There are many risks that can threaten the financial well-being of insurance companies, and today I'm going to give some examples found in the 10 largest U.S. life insurer insolvencies in the last 20 years.

First, I'll give a very brief overview of the U.S. regulatory process for identifying and administering insolvencies. "Insolvency" can be defined as the inadequacy of assets in any way, signaling the possibility of future defaults on liabilities, resulting in actions by a state insurance commissioner ranging from temporary conservatorship to liquidation. The key here really is action by the state insurance department.

As far as a brief overview of the U.S. regulatory process, the same process is used for reinsurers and direct life insurers. In 1991 there was an unusually large number of failures. It really alarmed the regulators, and so the result was an accreditation process for the states. The states had to agree to adopt these processes for identifying and administering the insolvencies. The first part of this process is a quarterly review by the state insurance departments of each company's financial filings. Then companies deemed possibly at risk are discussed monthly at the NAIC financial analysis working group meetings. The next step would be an on-site examination, and if the company is still considered at risk, a rehabilitation order and plan. Often there's a temporary freezing of assets to prevent a run on the bank. For direct insurers only there's involvement of the National Organization of Life and Health Guarantee Associations (NOLHGA). In each state there's a guarantee association, which is a nonprofit entity that pays benefits on insolvent companies, and to the extent there are not enough assets to fund the benefits, the organization assesses the healthy insurers in the state to get the money. Then the final outcome would be either conservatorship for the company, which would be where the company continues to be run by the state; rehabilitation, where the company eventually can stand on its own again in some capacity; or liquidation, where the company would cease to exist.

As far as reinsurers in the United States, I looked in a lot of places and I talked to a lot of people, but I could find no evidence of insolvent U.S. reinsurers. Apparently some reinsurers have had their problems, but they've managed to avoid insolvency by either recapture, retrocession, novation or sale of blocks of business. A current example of a troubled reinsurer, and this is not a U.S. company, is Annuity and Life Reassurance, a $2.3 billion holding company based in Bermuda. It had a lot of healthy business, but it had three very large treaties that were not healthy, and eventually two of them resulted in arbitration and one in recapture. It suffered significant operating losses, and as it started having financial problems, its letter of credit providers were demanding collateral. The company tried to raise money but was unable to get favorable terms. So it was downgraded by all the major rating agencies, and as of March of this year it had $172 million in collateral shortfall and ceased accepting new business. So that's my one reinsurer example.

I'm going to talk about the top 10 insolvent direct insurers, ranked in order of the year of insolvency. When I say "recovery of liabilities," this is meant to pertain to ordinary creditors. So 100 percent in this context doesn't necessarily mean that everybody was happy about the outcome. Another point is that it really takes quite a while to wind up these things in most cases; it is not a quick process.

The first company, Baldwin United Group, is not really a company. It's really six affiliated companies. There were three in Arkansas, the largest of which was National Investor's Life, and three in Indiana, the largest of which was University Life. The primary cause of failure for this group was that they sold a tremendous amount of single-premium deferred annuity business. Single-premium deferred annuities really require a fairly large amount of capital to fund acquisition costs and statutory reserves. So there was capital inadequacy, and this type of business particularly accentuated it.

They sold a lot of these contracts through stock brokers, and when this all went down, the stock brokers who had been selling this stuff got together and contributed money to try to help give the policyholders more interest. Also the investment industry and stock broker industry contributed money to help defray this and make the policyholders a little bit better off. The existing policyholders were eventually given a choice of cash surrender or a policy from Metropolitan Life.

The second company on the list is Mutual Benefit. This was a highly respected, old line mutual. This is my first example of the 1991 companies; there were 25 in all, and five of them made this list. Failure there was caused by commercial mortgage and real estate losses followed by escalating policyholder surrenders. The liabilities were restructured and transferred to a subsidiary of this company, MBL Life Insurance Corporation. The ownership of the stock was transferred to a trust, and the insurance commissioner was the sole trustee. They had quite a bit of group pension liabilities, and they were transferred to a separate account and guaranteed by a consortium of insurance companies. The policyholders who opted out of this received 55 percent of the available account value. Eighty-two percent in this context was kind of an average of the people who stayed in and the people who got out.

The next one is Executive Life. This very famous case may or may not be fully resolved. There are still things ongoing, which I'll tell you about in a minute. At Executive Life the failure was caused by defaults and investment losses on high-yield bonds. They also had a problem with policyholder surrenders. At year end 1990, which was the year before insolvency, 94 percent of the company's long-term bond holdings were noninvestment grade. A French consortium bought the assets and the liabilities and transferred them to Aurora Life, and this seemed like the end of it for a while. However, in 1998 the State of California, which was the state of domicile, got word that basically the French consortium was a few front companies for the French bank Credit Lyonnais, and, in fact, at that time it was illegal for a bank to own insurance companies in the United States. At the time Credit Lyonnais was a state-owned bank. So the State of California issued a lawsuit against the French government as well as Credit Lyonnais, and now they're in negotiations. This has been going on for over eight years, so it's a long one.

The next company I'm going to talk about is First Capital. First Capital's failure was also caused by defaults and losses on high-yield bonds, again with escalating policyholder surrenders. In 1991, which was right after they were taken over by the state, 29 percent of the invested assets were still noninvestment grade. Eventually all of the policyholders and creditors were paid in full. The policyholders were allowed to surrender, or they could continue with a restructured policy from Pacific Mutual.

The next company, Fidelity Bankers, was an affiliate company of First Capital. They had the same problem—high-yield bonds with increasing policyholder surrenders. That year, 1991, was a really bad year for that. The policyholders at Fidelity Bankers were guaranteed 100 percent account value if they continued as a contract holder with Hartford Life, or they could have 85 percent account value if they surrendered.

The next company is Monarch Life. This one is a different case. Failure here was really because of failure of the parent. The parent was Monarch Capital, a holding company, which had some real estate that it wrote down. I guess they had a lot of it, and when it was written down, it put them in financial trouble. Unfortunately they had some credit arrangements with banks, and Monarch Life was used as collateral in these arrangements. So when the parent was in trouble, the state stepped in and took Monarch Life to separate it from the problems of the parent. The parent was reorganized as Rico Reinsurance, and it's a life insurance holding company. It still is the parent of Monarch Life, which is to this day under the control of the commissioner in an effort to stabilize its financial situation.

The next company is Fidelity Mutual. In 1992 the failure there was primarily caused by poor real estate and mortgage loan investments. In 1990, which was about a year and a half before insolvency, 50 percent of the assets were real estate and mortgages. The rehabilitation plan is still being fully worked out. I think they're just about there, but the rehabilitation plan specifies full reimbursement to creditors, including 6 percent simple interest, which is pretty good as far as these guys go, as well as full values and benefits to policyholders. All of the business will be assumed by a subsidiary of Fidelity Mutual called Fidelity Life; of course, this is a mutual company. The policyholders will be given stock in the subsidiary in lieu of member voting rights. The conclusion of the plan and the transition to the stock company are expected to occur after 2004.

Next is Confederation Life, which is a Canadian company. It has a U.S. branch, Canadian Life Insurance & Annuity Corporation. My remarks pertain to the whole entity because the branch was really quite a bit smaller than the parent. Failure here was caused by poor performance of a large mortgage portfolio and the trust company subsidiary. At year end 1993, which is the year prior to insolvency, mortgages were 38 percent of the assets. The rehabilitators repackaged the mortgage securities into collateralized mortgage obligations in an effort to raise their value. Eventually the business was separated into pieces and sold off to different companies. This one is also still going on. Financial projections indicate that the ordinary creditors will be fully repaid with maybe 1–3 percent interest. However, subordinated debt holders will probably eventually receive between 55 and 75 percent. So this is one case in particular where recovery of liability is 100 percent plus interest, but not everybody was necessarily happy.

The next case is General American Life. This didn't happen very long ago. They had enough assets, but failure here was caused by lack of liquidity. General American had written quite a bit of short-term funding agreements, and they had reinsured $3.4 billion of it. Then they recaptured it from the reinsuring company, and that caused problems, leading the State of Missouri to step in. So in order to avoid a forced liquidation, the Missouri Department of Insurance basically took over the company and gave it a little bit of time to get things sorted out. Five months later General American was sold to Metropolitan, and Metropolitan completed funding the surrenders.

My last example is London Pacific Life & Annuity Company. This is a U.S. subsidiary of an English company. Its insolvency happened just last year, so this is very early in the process. Failure here was caused by losses from high-tech and telecommunications investments including WorldCom. Currently there's a moratorium on payments to creditors and policyholders, which is expected to last 24 to 36 months. The North Carolina Department of Insurance is in the process of getting a rehabilitation plan. If the investments can regain their value, it is expected that the liabilities could be fully funded.

There were 25 insolvency court actions that started in 1991. There were 12 in 1992, 11 in 1993, eight in 1994, two in 1995, four in 1996, four in 1997, four in 1998, seven in 1999, six in 2000, three in 2001 and only one in 2002. The number has come down quite a bit over the years. From looking at this I have two main conclusions. First of all, financially impaired U.S. reinsurers have been successful in using recapture, sale, novation or retrocession to avoid insolvency. My second conclusion is that the trend you just saw in direct insurer insolvencies reflects strengthened regulatory action and legislation.

MR. WASON: The next two speakers are going to talk about a development within the IAA with respect to solvency assessment. This project grew out of a desire by the IAA to maintain strong relations with the International Association of Insurance Supervisors (IAIS), the international body that represents insurance supervisors around the world. A couple of years ago they approached the IAA and asked if we might be able to assist them in thinking about a new solvency framework, so this working party that Dave and Harry are going to talk about was a result of that. We're just now poised with a draft of the final report, which is going to the committees and council of the IAA in November. Some of the importance and relevance of it to you is that a number of jurisdictions around the world are actually moving ahead to change their solvency standards. So I think this type of report is timely, and it's good to see the profession involved. I guess we'd encourage actuarial associations in many countries around the world to be active in this area and try to be out front and assisting the supervisors in this regard.

MR. DAVID SANDBERG: I'm first going to give you a little background on this working party. One of the early important decisions was that we would assume the American version of a committee or task force. This was formed in 2002, and basically the question that we're trying to frame for the group of international insurance regulators is whether there is a unified or consistent framework that they could start with as they begin to develop insurance regulation within their member countries. They wanted it to be both conceptually sound and practically focused. What is it we can do whether we're in the United States, Canada, Indonesia, Brazil or some of the other developing countries?

Certainly we were able to leverage off of a lot of the work that's been emerging in Canada, Australia and England. But we also want to make sure that it's workable for other people to get into the game of sound insurance regulation. In some countries the prime minister basically appoints his brother to head up an insurance company so they can funnel a bunch of money back into the family coffers. You need an objective framework for a country that shows how to make sure that the needs of the policyholders are being protected.

Last, and most important, the biggest hurdle to try to work through is to recognize that political and real-world constraints are very different. It's not just a case of picking up a model and running it. I feel that we've come up with some ways to try to address that issue. It's also important to remember that capital exists in a larger context than actuaries are used to thinking of, which is, "What's the number?" Or maybe even regulators, who say, "I have to set aside X dollars for my rainy day, and as long as that's there, then everything's going to be fine." We wanted to make sure that our approach looked at, discussed and emphasized the value of these four items, and that beyond the calculated number, there's an important need for corporate governance. In Canada they're starting to use the phrase that the first line of defense against insolvencies is sound corporate risk management, and the second line of defense is capital for the inevitable. We aren't going to think of everything, and we need to make sure we have some kind of rainy day fund, but the critical element is that the idea of ownership gets driven within the company. The company should feel like it owns its risk. You have to make sure your building has a sound risk-management structure. Certainly there is the question how you can include peer review.

Another thing that's important is the supervisory relationship. In a session early yesterday, Walt Rugland mentioned a kind of regulatory relationship in which regulators would make a rule, companies would spend all of their time trying to get as close to the line within the rule as they could, and then the regulator would make up more rules to make sure that he's catching those that are getting too close to the line. This is opposed to the banking relationship for large banks in which you often have the owner actually living within the company context. They get invited to all the meetings, they see all the memos, they have access to all the records, and, in fact, the management wants to be alerting them early on so that they have buy-in early on from the regulators.

We've touched on some of the keys, such as active company risk management and quality of regulatory oversight. You also want to make sure you have a minimum capital requirement. At some point you need the authority to come in and say that the company can no longer run on its own and management needs to be replaced. You also need to have some way of accessing guarantee funds for the bad cases that occur.

Our group was composed of a very diverse group of individuals from general insurance, life insurance and regulation. We are from Europe, the United States, Canada, Australia and Japan. Interestingly enough, we had a talk yesterday from our incoming president, who emphasized that one of the things that we need to be working on as a profession is our image. In order to show people that we're actually willing to stand behind our work, we have pictures associated with the people involved in the working party.

What are the key principles that are coming out of this document? First of all, we are doing this in the context that the international approach is to have a unified approach to both banking and insurance and that the Basel accord/guidelines have already spent a significant amount of time thinking through a multipillar approach. We need to look at all the rest of the company. We're not just looking at interest rate and asset risk. There is certainly value in having principles being espoused versus rules. One of the important elements is to try to avoid debates on international accounting standards and use a shortcut of an integrated balance sheet. As long as asset cash flows are linked to the liability cash flows in aggregate, they're sufficient, and I don't have to worry as much about whether I'm on a German GAAP basis or U.S. GAAP basis. I wanted to make sure we understood what's the appropriate risk measure. The definition of insolvency is when there's a ______ (fill in that blank) possibility that the liabilities will not be funded. So how do you decide what that possibility is, and what's the time horizon.

Another key element is making sure we get credit for risk dependencies. Everything will not go bad at once. At Allianz we actually have an internal risk aggregation model we've been prototyping worldwide. We do all the individual risks, add them up, and then there's a matrix that gives credit for dependencies. Typically we see roughly half of our total economic capital calculation being reduced because of the interdependency effects. So it's a significant element and certainly points back to making sure that management is owning the risk as opposed to renting the risk.

We talked about the multipillar approach. Essentially, if you think of pillar one as the financial calculations that are similar to U.S. risk-based capital, there's some dollar amount that's calculated. You have a second set of risks that are a little fuzzier. Maybe it's the risk that I don't know how policyholders really behave when interest rates are at 1 percent or 15 percent. Maybe it's an issue dealing with operation risk. You want to have some supervisory review and involvement and discussion of these items, and you also want to impose the discipline in pillar three of replying to certain disclosures. I think it's important to recognize that disclosure has context and is important in improving the balance sheet information. You want to be able to link to the balance sheet, but I don't think this means that you have a 100 page report that's disclosing additional information. But you do want to make sure that the company is taking responsibility for what it believes about the future and sharing that with the outside world.

So what are the challenges? We said it needs to be sufficiently simple to be applied anywhere and detailed enough to reflect companies' specific characteristics. It can't be one size fits all, and yet it can't be complex enough that everybody's using an internal remodel. Well, start with what the idea is and what we want to get to. We'd like to encourage the use of the internal model, identifying risk measures explicitly and building risk-based formulas with approximations to the internal model. In the United States we're just in the closing phases of finishing up the C-3, phase II project for capital for variable annuities. We've worked with this internal model approach, we've defined how it works and, as the last step, we've built some formulas or factors that companies can use. The problem is how do you develop factors that are broad enough? One solution is to use computers that work for three or four weeks. You have to run the models many times to capture all the ranges that you could need for the factoring.

The IAIS has a risk classification scheme that basically looks at asset risks, the technical risks, which are really the reserves, the traditional actuarial liability side of the balance sheet and then the nontechnical risks that we often refer to as operating kinds of risks. The banking system has its own risk classification scheme that is organized around credit risk, market risk and operational risk. We thought it would be helpful to propose another set of frameworks to introduce to the IAIS on how to think of the risk, which allows us to walk between bank classifications and insurance classifications. Underwriting risk, credit risk, market risk, operational risk, liquidity risk and event risk are the categories that we felt were important to identify.

I don't want to spend a lot of time on the specific risks, but I'll go through some of the details. I don't think there are surprises to anybody about the kinds of events that we're including in the risks.

Specific risks: The first three categories, underwriting, credit and market risks, are in pillar one. Under underwriting risk, there is underwriting process risk, pricing risk, product design risk, claims risk (for each peril), economic environment risk, net retention risk and policyholder behavior risk. Under credit risk there is business credit risk, invested asset credit risk, political risk, reinsurer risk and sovereign risk. Under market risk, there is interest rate risk, equity and property risk, currency risk, basis risk, reinvestment risk, concentration risk, asset-liability management risk and off-balance-sheet risk. One question you have is where to draw the line and say, "This isn't a risk requiring quantification, but it's a risk that we need to disclose and have dialogue with the regulators." The fourth risk, operational risk, sits on the fence. The banks have said we'd like to have an operational risk model, but they currently say 8 percent of your capital is defined as how much is needed for operational risk. We recognize within the insurance industry that it's even more difficult to come up with an operational risk model. Unlike banks, where their risks are transactional based and they're doing hundreds of thousands of transactions a day, insurance operational risk has more to do with fraud and systems failure. Under operational risk, there is human capital risk, management control risk, system risk and strategic risk.

The last two on the list, liquidity risk and event risk, are also two issues that are not under pillar one, but they are part of pillar two. What does pillar two involve? Another way to think about it is similar to the two lines of defense: capital risk management and the money set aside for a rainy day. Pillar two is really about the risk management process. Can a regulator come in and evaluate the process that the company's engaged in and then come up with some qualitative assessment of A, B or C or high, medium or low that gives them some feel for how well positioned the company is to address these kinds of risks? Under liquidity risk, there is liquidation value risk, affiliated company risk and capital market risk. Under event risk, there is legal risk, reputation risk, disaster risk, regulatory risk and political risk.

We've talked about integrated balance sheets, but which risk measures should we use? We have mean, standard deviation, value at risk (VAR), CTE, TailVaR, TVaR and other risk measures that you may use in your own company. We'll go through the 75th percentile, or we'll have some other measure that chooses the pricing mechanism. Our esteemed colleague Harry Panjer will discuss why we ended up recommending TailVar and the value that comes from that. CTE, TVaR and TailVaR are all part of the different ways that this has been described, but it is the same risk measure.

Time horizon: We propose two tests, a short-term test based on, let's say, 99.5 percent TailVaR, and a longer-term test, say, at 95 percent TailVaR, that's looking at the lifetime of the business. Given that we want the internal model to be used, in practice we want some simple methods that can allow them basically to use an exposure measure and a factor and then a formula to combine all the products that are there.

Now, one of the nice fun things that I think comes out of this paper is that we framed the use of what are called copulas. Basically it's a way to describe relationships and dependencies between different distributions, even when you don't know the exact relationship. So it allows you, as a regulator, to say, "I don't know what the correlation is, but let's say it's between this kind of balance. Let's put a range on it, and as long as we think it's in this range, we have a way to start aggregating and putting a number on that." It then allows you to start systematically pointing your research and analysis at getting more precise about estimating that relationship.

The key critique of the internal models is that data on tail events are difficult to capture, and it's difficult to get agreement on what may be occurring in the tail. Again that's why we have the event risk, and some of the operational risk issues being in the phase II requirements as opposed to part of phase I. Given that we went through the conceptual framework, we realized we really need to have some way to assess this. So the real work that I think was invaluable is that actuaries in these three lines of business made a fictitious company model, applied an internal model using the framework, and then came together to think about factors that you could use.

What was interesting to me is that on the property & casualty (P&C) side I think we ended up with a very workable framework for developing factors based on applying the theoretical approach that was used. We were not as successful in coming up with a generic answer on the life side, but we made significant progress. In fact, one of the ways that we've realized that this framework is going to be most successful is that this is not being sent out like the Basel approach. Basel I came out, and then banks adopted it around the world. The requirement was if you want to trade with another bank in another country, you must comply with Basel or there will be restrictions on what you can do. It led to an immediate, very uniform, adoption of that standard throughout banks around the world.

We don't have that ability in the insurance environment, nor are we looking for that kind of approach. What we want is the framework. We finalize this at the IAA meeting in November, get feedback from different countries and then put out a final report out next spring. We're basically going to turn this over to the national actuarial organizations as their template for dialoguing and working with the regulators in their country. So that is the framework and the key principles, and you can use it to figure out what it means for your country. You will be able to ask where you fit and then let the local actuarial organizations work with the regulators in order to get the degree of regulatory oversight, actuarial involvement and peer review that makes sense within both the industry and legal traditions of your country.

MR. HARRY PANJER: I'm going to talk a little bit about some of the more technical aspects of the work we've done. Even the technical work is at a relatively early stage. There is another working group behind this working group. It's really a research group in which I participate. It is doing some of the technical development, and we're discovering some interesting things that I think could contribute nicely to the development of these formulas.

I'm going to talk about three aspects. One is risk measures and how we measure risk. Then I'll talk about dependency and how we model and measure dependencies. The third part will be developing factor-based methods from internal models.

The hypothesis under which we are working is that a well-developed internal model is the best framework for assessing risk for an individual company. The first degree of simplification is to develop some factor-based methods that reflect the essential characteristics of that model. So we try to pick up most of the characteristics of the model through some kind of approximation, and there are different ways of developing approximations. We've investigated a couple, and I'll talk about those.

First, on the issue of risk measures, classically and even currently in finance, most people think in terms of value at risk. Value at risk is another way of expressing what we've always done in actuarial science, which is probability of ruin. What's the probability that you're ruined with this amount of capital? The value at risk question is how much capital you need for a fixed probability of ruin, so it's really the inverse question.

There are other possibilities. One of the key characteristics of risk measures that has been discovered recently and seems to provide the properties that one would like in a risk measure is what is called coherent risk measures. This is work that has come out of the finance literature and was done primarily by actuaries in Europe. Coherent risk measures have certain sensible properties. The first property is being subadditive. Capital for two risks is not larger than for each risk separately. The second is that risk with no uncertainty requires no capital. The third is invariance under location and scale transformations, such as changing currencies. The last property is additivity for comonotonic risks. I won't go through them in detail, but value at risk and the probability of ruin are not coherent. They have some pathological problems when you try to apply them in situations where you're aggregating or disaggregating risks. You don't get consistent results or logical results when you use value of risk on its own.

One interesting property of coherent risk measures is that a coherent risk measure expresses the capital as an expected value. Now, the capital requirement here is the total balance sheet requirement. So it expresses an expected value of a distorted distribution (Figure 1). It's distorted in the integral by the distorted function *g*, and *g* is just a mapping from the unit line to the unit line, from the unit interval 0, 1 to the unit interval 0, 1.

Figure 1

Coherent risk measures

- Capital requirement can be expressed as an expectation under a "distorted" measure

$$C = \int x \cdot dg(F_X(x))$$

where $g(x)$ is a concave continuous function on the unit square with $g(0) = 0$ and $g(1) = 1$.

- Every coherent risk measure is characterized by a distortion function.

University of Waterloo

1/6/2004 SOA PD126 October 29 2003 5

Now, what is important about this is not this mathematical form, but to recognize that actuaries have done this all along. We've always done things like adding margins. When we add margins to probabilities, then all we're doing is changing the probabilities and, therefore, changing the underlying probability structure. That's all this does, so it's essentially consistent with that notion.

TailVaR, TVaR or CTE is shown in Figure 2. It's a risk measure that says if you have some quantile x_q, what you want to do is represent the total balance sheet requirement as the expected value of *x* conditional on *x* being greater than that amount. So the idea is you might pick, say, the 95th percentile of the distribution. If you think of the outcomes as a distribution, then you pick the 95th percentile and say we'll condition on the bad event occurring. We're going to have an outcome we know is larger than that 95th percent quantile. We'll take the average value of those events that exceed that quantile, and that will be our total balance sheet requirement. From that we will subtract our reserves, and the difference will be our required capital.

Figure 2

Recommendation

- TailVaR (CTE, TVaR) as risk measure
  - Find x_q satisfying

$$\Pr\{X > x_q\} = 1 - q$$

where X represents loss to the insurer.

  - Total balance sheet requirement (reserves+capital) is

$$\mathrm{E}\left[X \middle| X > x_q\right]$$

University of Waterloo

1/6/2004 SOA PD126 October 29 2003 6

I know that's already in use in Canada with respect to certain products as of a year or so ago. It's called a shortfall constraint because there are really two parts to a tail value at risk. One is a value at risk that is the 95th percentile, the amount you need for a 95 percent probability of security. The other is the expected shortfall, which is the overshoot in probabilistic terms over the expected shortfall if you know that you are above that boundary. One of the nice properties of this is that it picks up part of some characteristic of the tail of the distribution, and it picks up the mean only. It picks up the average amount by which you will exceed that quantile, given that you exceed it. That's kind of first-order information about the tail.

Now, why is this important? Well, when we look through the range of insurance products, from property, casualty, catastrophe to life, we have different shapes of tails of the distributions. P&C actuaries like to work with the Pareto distribution because it's quite flat in the tail. It's typically a group life insurance portfolio that doesn't have any significant investment component and has a very light tail, much like a normal distribution. We would say for those two, even though the value at risk may be the same for two different product types, that we would expect to have more capital for the one with the heavier tail. If you're going to have bad luck, you're going to have really bad luck in the P&C case compared to the life insurance example I mentioned.

In actuarial terms the expected shortfall is actually the part of the stop-loss premium that is conditioned on that stop-loss claim occurring, so we can describe it in terms of stop-loss premiums. For a normal distribution, the tail value at risk is larger than the value of risk, of course, by that average overshoot, and you see that in Figure 3. This difference is the excess amount. Now, of course, the tail value at risk is itself a value at risk because it's just a quantity along the X axis; value at risk is one quantity on the X axis. This is just a higher percentile. The point is, for different distributions, this difference will be different.

Figure 3

TailVaR

$$
\begin{aligned}
\text{TailVaR} &= \mathrm{E}\left[X \middle| X > x_q\right] \\
&= x_q + \mathrm{E}\left[X - x_q \middle| X > x_q\right] \\
&= \text{VaR} + \text{expected "shortfall"}
\end{aligned}
$$

- Expected ***shortfall*** is the ***net stop loss premium*** for excess losses given that a stop-loss claim occurs.
- The ***trigger point*** x_q can be thought of as the point at which the current assets are just sufficient (on average) to cover current liabilities.

University of Waterloo

1/6/2004 SOA PD126 October 29 2003 7

Actually a tail value at risk could also be considered the average value at risk, as is shown in the shaded area of Charts 1 and 2. If you could choose all the quantiles and just average the value at risks, you will get that tail value at risk. So it's another way of thinking about the tail value at risk. In this case it's the average of all values at risk with the trigger point being the percentiles in excess of 95 percent.

The real reason we focus on tail value at risk rather than value at risk is that different insurance portfolios have different risk characteristics, and that translates directly into having different distributions. There are other at risk measures that you can come up with, but this one seems to be a very sensible thing, and it's relatively easy to work with. It's easy to simulate, because what you do is, you simulate, and you take your worst set of outcomes, and you just average them and you say that's the total capital requirement. In this case if you want the 99.9th percent level with 10,000 simulations, you just take the 10 top reserve values and average them, and that's the total required capital including reserves.

The last point is very significant compared to value at risk. Value at risk requires that you come up with a quantile, and if you're doing a simulation in the extreme of a distribution, and the distribution itself is quite flat like a Pareto distribution, you can end up with observations that are quite spread out when you've got very few observations in the tail. There are different rules for coming up with what the quantile is in the 95th percent quantile or the 99.9th that we have here, because you often have to interpolate between two values. That value that you come up with for value at risk is highly sensitive to the two observations around it, whereas the tail value at risk is much more stable. It depends on all of the observations in excess of a quantity, so it's a much more stable statistic when you apply simulation. That's all I'm going to say about the choice that we're recommending for a risk measure. Until we have a better theory developed, this seems to be the best that actuarial science or finance gives us.

Now, the question of dependency is an interesting one. We know that for most companies, if they're interested in modeling risk, it's done on what we would call in statistics a marginal basis. We can gather data on various types of insurance policies either within a company or between companies and put them together in a nice data file and come up with some distribution (Figure 5).

Figure 5

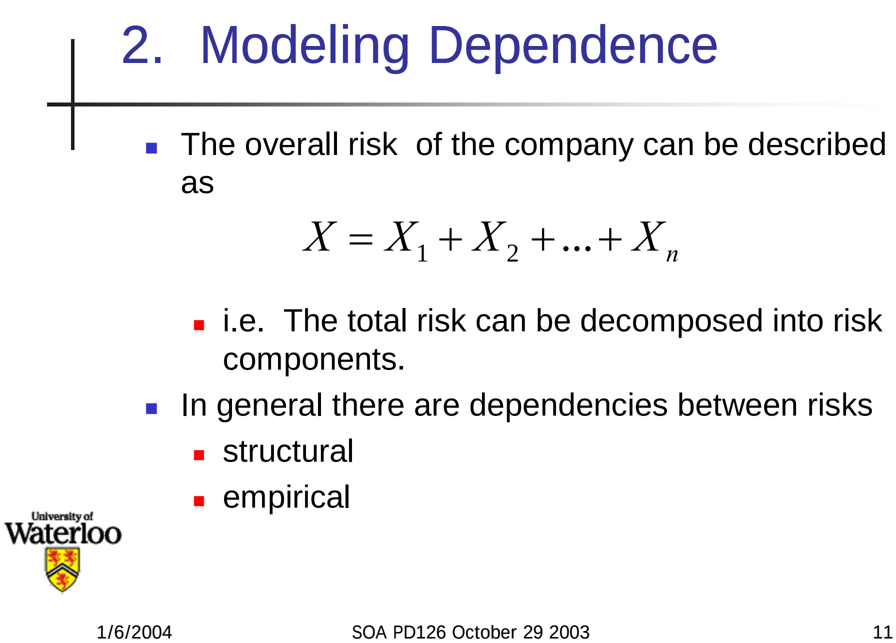

We know a whole lot less about the interdependency between outcomes for different risks. There are different ways of describing dependencies. One we call structural, and one we call empirical. The structural one is where we understand what the underlying drivers are. These are things like economic factors, such as inflation driving various lines of insurance, or common shocks, such as a single event triggering claims for different lines of insurance, the simplest being an automobile accident with claims for P&C insurance. An automobile accident can result in several claims for the different lines and different coverages within the policy. They may be claims on some but not all.

There are also uncertain risk variables. In the terminology of the Canadian actuaries, it's the uncertainty about the mean. In other words, how well do we know the statistical uncertainty about what the average value is. That's kind of a systematic error, in that it is a risk that cuts across all policies. If you get the mortality rates wrong, you don't get them wrong for one policy, you get them right or wrong for the whole block of business. That is a structural dependency. If we know the nature of that uncertainty based on, say, standard errors in a statistical analysis, we can come up with some estimates of what that uncertainty is. Then we can incorporate it into the model, and so it's structural. We can build that structure into the internal model.

There are, of course, catastrophes. A catastrophe is really a common shock, but it's an extreme common shock. We've identified event risk as a special risk that we call an extreme event risk. We want to make sure we recognize extreme event risk in the insurance business, particularly the P&C, the reinsurance or the P&C reinsurance areas. They're typically not recognized in the Basel-type structure for risk recognition within the banking system. Catastrophes can affect many lines of insurance and can affect other risks as well, such as the company's own operations, so it also could be an operational risk.

The point of structural dependencies is that known relationships can be built into internal models. Now, more typically we don't know what those relationships are in a structural sense. We don't know the cause-effect relationship, but we do know that things can go bad all at once. We typically don't have very much data on relationships because we're very much focused on extreme events. If we're interested in insurance company solvency issues, we're interested in the extreme outcomes. We're not so much interested in the dependency that occurs on a normal day-to-day basis when things are operating within the normal range. We're especially interested in the cases where things really go wrong, and we haven't much data.

What we can do is build models that hypothesize a certain relationship, and we can do that mathematically. The mathematical tool that we will use is the copula. Chart 3 is a set of outcomes. These are purely simulated data for two variables, and these are on the 0, 1 axis. We can think of these outcomes as being just the smallest to largest, so we just rescale to get the real outcomes. You notice that in this picture there is some dependency. If you look at the middle of the picture, it looks kind of normal; it looks uniform. Once you start to look at the extremes, then things are more clustered in the lower left than the upper right. Let's think of the upper right as being the bad events and the lower left as the good events. We're most interested in that upper right point. For example, if you look along the right-hand side of this square, you'll see that if one of the outcomes is bad, there's a tendency for a bad outcome with the other variable as well. For this purpose we're much more interested in this upper right corner than we are in the rest of the picture. So in normal times we might see what's in the middle, but we're interested in building models that allow for the kind of concentration of bad outcomes that is exhibited in this diagram.

Here's how copulas work, and let me tell you that copulas are nothing new. A copula is just a way of describing a multivariate distribution. Copulas were discovered in 1953 but really ignored for at least 45 years. First of all, let me go back one step and say a copula is a way of describing a multivariate distribution in terms of the marginal distributions. Let me explain what that means. If you think of a multivariate distribution, you could think of just a bivariate distribution. You could just picture it in two dimensions, some sort of hill for the different shapes. You could think of that hill in terms of three components: the projection in one direction, the projection in the other direction and then a function that links those two hills. The copula is a function that links those. That's all it is, and it does it by linking the ordering of the outcomes of the variables along the two axes. So it does it in terms of relating the magnitudes—magnitude not in an absolute sense, but magnitude in terms of the order or the order statistic. So it links large outcomes with large outcomes.

When it was developed in 1953 it was called a Scar serum, and the technical result is that the linking through the ordering of those risks is independent of the marginals. Therefore, when you think about how you put this together, you can separate out the data you have on the marginal distributions for different lines of business. This is what I expect from the dependency model that links them. That's a really powerful result because we know we can get data. We do intercompany studies for mortality. Reinsurers collect data and build models for certain types of risks. The copula is a way of linking them together, and it's independent of the shapes of the distributions of the two risks that you're putting together.

The favorite copula at the moment is the T copula. There are a number of other copulas as well, but it was a T copula that actually generated these data, and it's really easy to work with. If you can generate normal random variables, then you can generate T copulas. The T copula comes from the T distribution, and you can control very explicitly with one parameter the degree of dependency. There have been a lot of studies for different types of risks in banking, some in insurance, some in stock market returns and so on. With all of these the typical degrees of freedom are between three and five, so there's a relatively narrow range of degrees of freedom that has been observed in the marketplace. We have some guidance on what degree of dependency we might use in a real application prior to coming up with detailed studies for exactly the kinds of risks we have at hand. So the T copula allows us to reduce the dependency to a single number. That number, of course, is independent of the choice of marginal distributions. If you have a P&C line of business, which may have a Pareto distribution, and some other line of business that has a normal distribution, you can put them together with the same confidence that you can any pair of distributions.

This is all very recent research, and there's a lot going on. Other copulas are being developed, some of which have dependency only in one corner, so the lower left corner does not observe dependency, but all they are mathematical functions that relate the marginal distributions to a multivariate distribution. That's really quite technical, but actually in practice, once you get accustomed to thinking this way, it's quite easy to think in terms of a single dimension, and that's a degree of dependency somewhere between three and five. Where this dependency modeling using copulas is very useful is actually in stress testing, and it's the stress testing of dependency itself. In other words, you can ask what happens if we have a higher degree of dependency, and we can reply that your capital requirements should be larger as a result. We don't know what it is, but if you test it at levels three, four and five, then you'd get an idea of what the sensitivity to the capital requirement is of dependency itself. So dependency then becomes one of the variables upon which you can test.

My next topic is factor-based methods. The idea for developing factor-based models is that you start with an internal model that captures all aspects of the risks and their interactions. We assume it's the best model. Now, that's a hypothetical model because none of us have that perfect model, but we start with the hypothesis that the perfect model gives us all the information that we need. What we need to do is extract as much information out of that model as we can in terms of the key characteristics and turn it into a simple formula consisting of exposure measures, a factor to apply to each exposure measure and a formula to combine all of the products. That's essentially U.S. risk-based capital (RBC). That's what we're looking for, that kind of simple formula.

The reason for doing this is to develop a simple formula we can use in any jurisdiction. The ideal framework would capture all key characteristics of a company's risk, including all sources of risk under pillar one and all interactions between different risks. It could be used in any country around the world at different levels of sophistication. We can combine, for instance, lines of business or aggregate risks in a formula like in U.S. RBC, or we can do this aggregation at higher or lower levels. Our challenge is to come up with a formula that will allow a company or allow a range of countries to do this at different levels of aggregation of data and different levels of sophistication, but have results that are internally consistent. Actually U.S. RBC is kind of the ideal reference point for doing this kind of work.

However, it requires company-specific calibration. Data on extreme events are very thin. You heard already from Liz about the failures of a number of life insurance companies. It was quite interesting that there were not many risks involved in the failures of those. One that we typically call concentration risk, which is when a big chunk of your assets is in a single category and it's going bad, came up. Liquidity risk came up in a couple, and one that was in our classification of risk, affiliated company risk, did come up. There weren't too many different risks that have actually caused failure in the system in recent years.

Building internal models also typically requires extensive model development and data collection. This is something that large companies and sophisticated countries can do. We can do it in Canada or the United States for most large companies. For small companies, particularly P&C companies, and small developing economies, developing an internal model is going to be something that they can't handle. Either they don't have the capability for doing it, or they don't have the people skills within the company to actually build it or analyze it. Sometimes it may not be necessary for smaller organizations. If we have rough rules of thumb where we know we have plenty of capital, there's no point in building an internal model just to validate that fact.

Calibration is difficult in the tail, especially within a single company. A formula approach may actually be a bit more useful in the extreme tail where we have little data because we can depend on more studies like Liz gave here, which are based on intercompany data.

The U.S. RBC formula looks like this, where we have a formula that brings together the indicated capital requirements for different risk types. These risk types include the following and their formulaic designations: $C0$ is affiliated investments, $C1$ is asset default risk, $C1_{cs}$ is unaffiliated common stock risk, $C1_o$ is all other $C1$ risk, $C2$ is insurance risk, $C3_a$ is interest rate risk, $C3_b$ is health credit risk and $C4$ is business risk. Each category has many risk elements, and each risk element involves product of exposure measure and a specified factor (Figure 6). The U.S. RBC covariance adjustment formula that is used with these is as follows:

Figure 6

US RBC "covariance" adjustment

- RBC = $C0 + C4_a + \sqrt{C1_{cs}^2 + (C1_o + C3_a)^2 + C2^2 + C3_b^2 + C4_b^2}$
  - Recognizes likelihood that not all risks will occur at the same time; i.e. lack of correlation of some risks
  - Uses correlations of either 0 or 1 for simplicity
  - Exact if standard deviation is a risk measure and correlations are correct
- However
  - Insurance company risk is often not Normal
  - Better risk measures should be used to reflect tail risk

1/6/2004 SOA PD126 October 29 2003 21

This formula recognizes the likelihood that not all risks will occur at the same time, that is, the lack of correlation of some risks. Implicit in this formula are correlations of either zero or one between risks. It uses these correlations for simplicity. Think of the risk measure as being the standard deviation, because this is how standard deviations are added together. So when we have two quantities that we square and add together, we end up with a square of a new quantity. That's what we do with standard deviations: We add them for independent random variables. So wherever we have squares of quantities added together, we treat those two quantities as having correlation zero. When they're added together with squaring as is indicated in the two in the fences there, we're assuming that those two are perfectly correlated. In other words, that's kind of a worst case.

What are the efficiencies in this one? We could be more accurate about the zeros and ones. I'm assuming your correlation of one is very conservative, because I'm assuming the worse case. The second is using a standard deviation as a risk measure. We've already recommended something more sophisticated, and that is tail value at risk, which is a better indication of the risk in the extreme of a distribution, whereas a standard deviation is measured over the whole of the distribution and really concentrates in the middle.

Figure 7 shows us what we do in terms of formula development. We say the total balance sheet (TBS) requirement, which we generally refer to as capital but is really capital plus reserves, is the expected value of *X*, our outcome variable conditional on *X* being larger than some quantile. This is like the reserves plus CTE or TailVaR. What we do is break this TBS up into two parts. We can say it's the mean of the distribution. This would be your outcome on a best estimate basis. Stick best estimate factors in for everything, look at the outcomes and look at the average outcome, and μ is that average outcome. It's the reserve that you would set up on a best estimate basis without adding any margins whatsoever. The *k* multiplied by σ part is a capital portion, but it includes all of that systematic risk that's incorporated into the safety margins and reserves. So it's safety margins and reserves plus what we now call pure capital.

Figure 7

Development of formulas

- For an internal model, total balance sheet requirement is

$$TBS = E\left[X \middle| X > x_q\right]$$

- This can always be written as $TBS = \mu + k\sigma$.
- The "capital" is obtained as

$$C = TBS - \mu = k\sigma$$

For Normal risks, the value of k can be calculated easily.

- For an entire company the distribution is likely not close to Normal, so more detailed analysis is required; e.g. heavy tailed distributions will have larger values of k .

University of Waterloo

1/6/2004 SOA PD126 October 29 2003 22

Sometimes those are described in terms of systematic risk and nonsystematic risk, but those definitions aren't particularly valid in all cases. Sometimes a systematic risk can be a nonsystematic risk when it's combined with other risks. So the pure capital portion is *k* multiplied by σ. Now, σ is the standard deviation of the distribution. The capital is just some multiple of it. All that does is take us along to some point in the distribution, and we just express that point as a multiple of the standard deviation of that distribution. If it's a normal risk, we can just look up in tables what this *k* should be. You can calculate this quite easily. For an entire company, the distribution is likely not to be close to normal, so more detailed analysis is required. We have to consider heavy tail distributions, and therefore we'll have a higher multiple of the standard deviation. So if you take two distributions that have the same mean and the same standard deviation, the one that has a heavier tail should result in a *k* that's larger, because you want to have more safety margin.

How do we put these together? The obvious formula is the kind of formula that we get from normal distributions, which is essentially the thing that's behind U.S. RBC now. You end up with a formula that looks like the one at the bottom of Figure 8. U.S. RBC is based on this idea. Risks are pulled apart, and correlations are put in at a slightly different level, but this is the idea. The C_i's are kind of the line-of-business capital requirements. This could be done at any level. The C_i's and the C_j's in the formula at the bottom have a kind of correlation, but it's not necessarily the usual correlation. For a normal distribution, it is the usual linear correlation. We could have a correlation that's based on other things, and we'll look at that in a second.

Figure 8

Adding dependencies

- Models are developed for specific risks within lines of business (LOB) and combined, resulting in

$$C_j = TBS_j - \mu_j = k_j \sigma_j$$

- LOBs are combines recognizing the dependence between them. So some kind of "correlation" is needed, say, $\rho_{i,j}$

- This suggests the simple formula

$$C = \sqrt{\sum_{i,j} C_i C_j \rho_{i,j}}$$

University of Waterloo

1/6/2004 SOA PD126 October 29 2003 23

The first formula on Figure 9 looks complicated because you have a mean, a standard deviation and a *k*, and you need all of those. When we talked to the international supervisors, they said don't make it any more complicated than the second moments. We've expressed the standard deviation as a multiple of the mean. That actually can simplify the formula so it does not involve the standard deviation directly, but a coefficient of variation as an indication of the shape of the distribution. Then we simplify and reduce the mean to an exposure measure and risks, as shown in Figure 10. For all of these elements we want to make sure the data are as simple as we might get now, so only *e* would come from the company itself. A regulator would determine all of the other factors based on intercompany studies, assuming, of course, that similar lines of business would have similar results for different companies providing that similar coverage. We've tried to take this and reduce it to a level in which the data input by the company is still very simple. It's exposure measures just as we have now.

Figure 9

Another representation

$$C_j = k_j \sigma_j = k_j \mu_j \nu_j$$

where ν_j represents the "coefficient of variation".

- The expected loss can be written as the product of an exposure amount and a standard "risk per unit".

$$\mu_j = e_j r_j$$

1/6/2004 SOA PD126 October 29 2003 24

Figure 10

Sources of data

- ν_j depend on shape of distribution, and is similar for similar risks for all companies, so this could be based on industry data.
- k_j depends on shape of the distribution and risk appetite of regulator. It is also then similar for all companies.
- r_j is expected loss per unit of risk and depends on industry data.
- e_j is exposure base and depends on company data.

- The "correlations" reflect risk measure, and copula or other measure of correspondence and so can be set by regulator.

1/6/2004 SOA PD126 October 29 2003 25

Chart 1

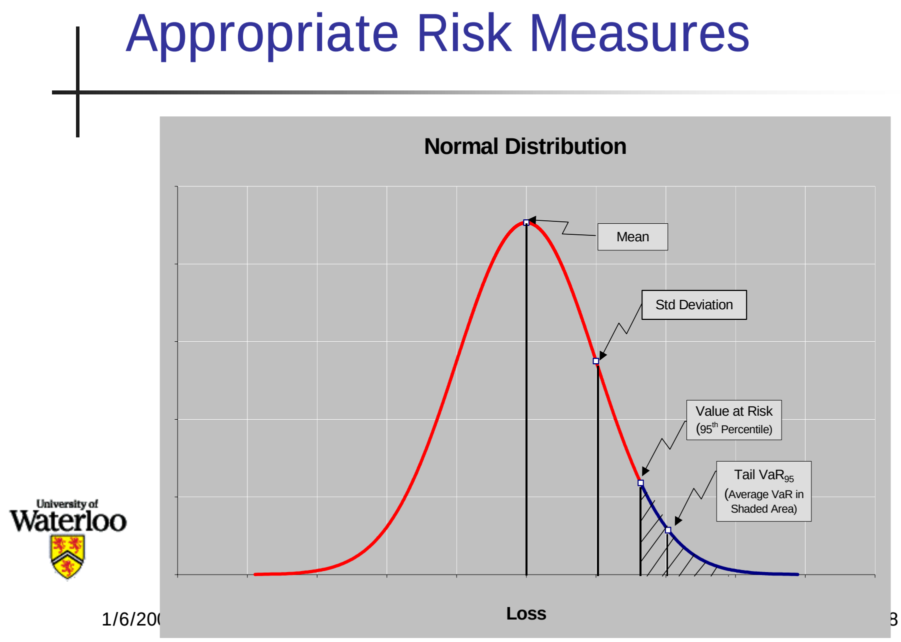

Chart 2

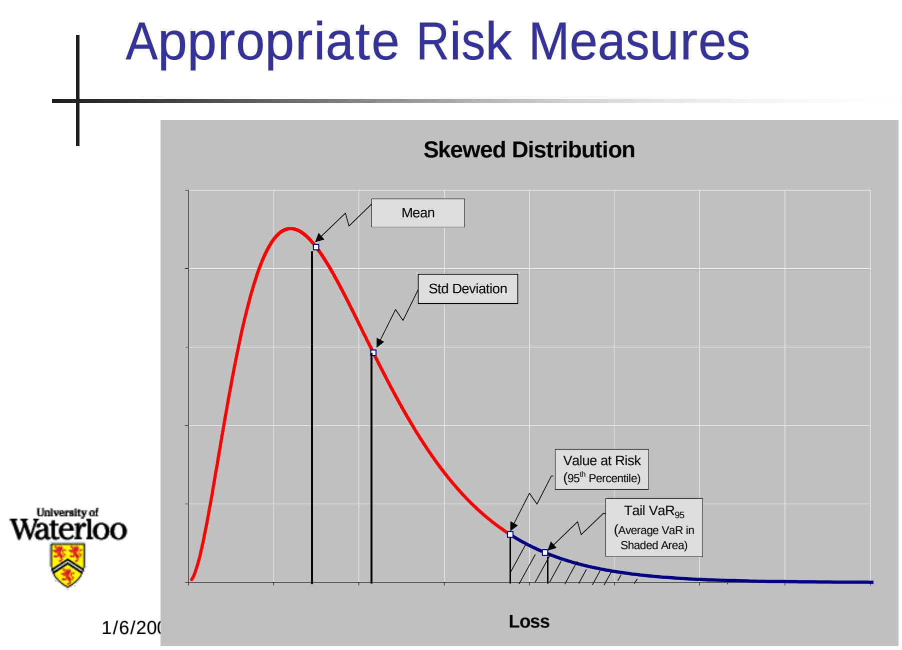

Chart 3

t Copula results

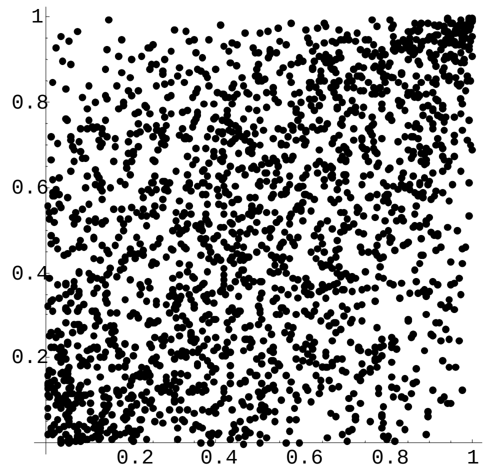

University of Waterloo

1/6/2004 SOA PD126 October 29 2003 17