

A review of the use of complex systems applied to risk appetite and emerging risks in ERM Practice

Abstract of the Edinburgh Discussion

[Institute and Faculty of Actuaries, 17 September 2012]

Contact

Neil Cante, F.I.A. E-mail: Neil.Cante@milliman.com

This abstract relates to the following paper: Allan, N., Cante, N., Godfrey, P. and Yin, Y. A review of the use of complex systems applied to risk appetite and emerging risks in ERM practice. *British Actuarial Journal*, doi: 10.1017/S135732171200030X

The Chairman (Mr D. C. Wilson, F.I.A.): I welcome our invited guests for this evening, Mr John Fraser and Mr Ian Campbell from the Registers of Scotland.

The paper was written a year ago in response to the profession's call for research into the areas of risk appetite, and the identification and assessment of hard to define risks.

Mr Wilson introduced one of the authors, Mr Neil Cante, F.I.A. Mr Cante also introduced the paper when it was discussed in London on 28 November 2011 and his introduction is recorded with that discussion.

Mr S. M. Gray, F.F.A. (opening the discussion): Risk appetite setting and understanding the emergence of risk are important. We are in the business of choosing to take on risk, so that we can provide customers with valuable products and investors with stable returns worthy of the exposures taken. A valuable by-product is that we also help to build sustainable businesses. None of this is new, so why the increased focus on risk appetite setting and emerging risk analysis? Our risk exposures are becoming increasingly complex and interconnected. Some of the interconnectedness has existed for a long time and was not evident but much of it is genuinely new.

Using a few examples:

- The impact of the US sub-prime mortgage market is now well understood: the dependencies within the financial system developed over time yet we were unprepared for the crystallisation of the risk event.
- The 2011 Japanese earthquake itself caused a lot of damage but the Tsunami it triggered had a massive impact. This illustrated the pre-existing interconnectedness of supply chains that had developed over time as a result of the proliferation of international trade. Again, companies were ill-prepared for the predictable impact of a break in supply chains.
- Why were we unable to deal with the impact of dust in the atmosphere as a result of the Icelandic volcanic eruption in 2010? An eruption was expected, and when it came it was relatively minor but the economic impact was huge. Once again, all of the information required to predict the

impact was available but our analysis and our preparedness fell short of what should have been reasonably straightforward.

If we step back from the detail of these examples we can see a powerful trend in the increasing interconnectedness which is changing the profile of our risk exposures and adversely affecting the diversification of our risk portfolios. It is for this reason that we must continue to investigate new techniques to help us manage risk exposures effectively on behalf of our stakeholders. With this driver in mind I welcome the paper that Mr Cantle has co-authored and I am excited to see how science from other domains can help us to make better use of the information that is already available to us.

The authors have articulated the problem in a framework. In this instance they have used a radar illustration to describe the necessary characteristics of methodologies to be used to address the risk management objectives. This is the first area where I would bring some challenge to the approach adopted. Based on my own experience of working with senior executives and non-executive board members, risk appetite setting and the analysis of emerging risk are already perceived to be very complex. Any new tool to assist needs to be intuitive and capable of simple explanation. Therefore my advice would be to add explicit criteria to the selection radar to reflect the need for intuitive interpretation. This would also help the techniques gain traction within the Actuarial Profession.

The second area of challenge relates to the method's capability to model the adaptive nature of the risks we are exposed to. Without this the methods will be overly sensitive to the starting conditions. It will therefore be important for the future development of complex adaptive science to take account of feedback loops and exposure to discrete discontinuities in financial systems, such as those introduced through regulatory change and corporate behaviours. That said, surprisingly valuable insights can emerge from innovative analysis of well known problems. As an undergraduate I gained exposure to *Fractal Geometry*. At the time I was fascinated by the beauty of the infinitely self-replicating graphs that could be produced by Fractals. Other than the production of the dramatic landscape of *Planet Genesis* in *Star Trek II*, I could not see any practical application of this branch of geometry. However, Fractals are being used today in medicine to identify heart defects and the growth of tumours. Fractals follow very simple rules and can produce outcomes which are highly complex and diverse. Who knows, perhaps Chaos Theory has a role to play in modelling emerging risks. The brief feedback in the paper from the case studies suggest that the methods in their current form have successfully augmented existing approaches in a life company to the setting of appetites and analysis of emerging risks.

My conclusion from this is that these methods have been able to extract new insight from existing data which is an encouraging result for techniques that are at an early stage of their application to risk management.

In summary, this is an important area for our profession to invest in. Success will bring substantial benefits to our customers and our business owners. It is particularly important to make progress, because the increasing interconnectedness of risks that we manage is changing the profile of diversification across our risk portfolios. I encourage you all to embrace the science of *complex adaptive systems* to help our profession be better equipped to predict and deal with uncertainty. It will not be easy, after 40 years and trillions of pounds of investment, meteorologists have consistently failed to predict accurately short term outcomes in the context of a slowly adapting physically observable system. But then again, if it was easy it would not be interesting.

The Chair: I now open the discussion to the floor and introduce Mr Neil Allan, one of the other authors of the paper. He is a fellow of the systems centre at the University of Bristol.

Mr C. G. Lewin, F.I.A.: For the last 15 years I have led the Joint Risk Management Initiative between the Actuarial Profession and the Civil Engineering Profession. Two of the four authors have also been involved in this initiative and have made extremely valuable contributions.

Starting with risk appetite, which is very much the subject of this paper. There are two concepts which have tended to be bundled in this paper. I see risk tolerance as different from risk appetite. If we take risk tolerance, it is really the toleration of ruin. Stakeholders might have a different risk tolerance from the board. For example, the board of an insurance company may be quite happy to hold sufficient capital to withstand all except a one in 200 chance of ruin in the course of the year. Policyholders' risk tolerance may be much less. A one in 200 chance of ruin means that there is a 10% chance that ruin will occur at some point in the next 21 years, which the holders of whole life policies may see as entirely unacceptable. That is risk tolerance, the toleration of ruin and the need to try to balance the views of different stakeholders on what risk tolerance they are prepared to accept.

Risk appetite is quite different. It is the appetite of the organisation for taking risk. A good example of a case where risk appetite vastly exceeded risk tolerance was the Halifax Bank of Scotland, where they had an enormous risk appetite for new business, which no doubt was based on the assumption that current financial conditions would continue, but the organisation was ruined when the financial crisis struck in 2008. Their appetite exceeded their capacity to bear risk. I believe that ERM would be clearer if we used both terms in our analysis rather than lumping them together as risk appetite, as the paper tends to do.

Taking another technique that is mentioned in the paper, concept mapping; this is likely to be an extremely useful tool in analysing and communicating risk. Just preparing the map forces a concentration on the underlying causes of risk, the connections between risks and the impacts that they may have on the organisation. Once the map is prepared, it shows which underlying causes are at the heart of these possible impacts, facilitating a focus on the causes which may best repay mitigation. So the concept map itself, leaving aside things like Bayesian networks, and so on, can be extremely useful in risk analysis. The authors then go on to recommend a Bayesian network approach to analysing the concept map, using computer software packages because of the complexities of the arithmetic.

First I tried to construct a very simple example of where Bayesian networks might be useful, building on an example in the paper. Suppose, for example, you have a new factory process and you are looking at the risks, and as a manager, have been told that all the staff have been fully trained. You reckon that there is an 80% chance that that is in fact true, and a 20% chance that they have not all been fully trained. You can assign estimated probabilities of having, say, five accidents a month if they have not been fully trained (say 40%) and another probability (say 5%) if they have been fully trained. And in the first month you do in fact get five accidents. Therefore what is now the chance that their staff were not in fact fully trained? If my arithmetic is correct, your 20% chance that they were not trained would go up to about 66% – it is just a very simple application of Bayes's theorem. Such examples, despite their simplicity, help in explaining the use of Bayesian concepts to people who might not otherwise accept that they are useful. Just the elementary arithmetic is worth explaining, because that helps boards to accept that the technique

could be valuable. In my example staff should probably now be retrained in order to reduce the risk in the future. So it would be a very practical kind of application. What the authors are doing is taking that much further and saying that they recommend using Bayesian computer packages in more complex situations. I am not sure how much confidence boards are going to place in the results from computer packages, given that they will be based on many different estimated probabilities. It is worth bearing in mind the lessons that were learnt when stochastic-modelling black boxes were applied to capital project appraisal. I am going back perhaps 15 years now. Merchant banks typically used the black box computer programs, which always produced bell-shaped curves for the distribution of outcomes.

But the problem emerged that nobody understood what was going on inside the black box. What were the assumptions? Were the assumptions right in a particular situation, and so on? It was just too complex for everybody. What has happened since is that the principal appraisal analyses have tended to be based on manual methods which everybody can understand. The black boxes are just used for confirmation, or bringing out certain technical points, but only as an aid rather than the principal method.

One of the fascinating aspects of the paper is the use of anthropological techniques to show how risks have changed and how they might change in the future. As actuaries, we should be historians. If you are looking at, shall we say, the trend of interest rates, you should always look back at how these have changed and why, before trying to make any judgements about possible future trends.

Exactly the same is true of risks. There are long histories of risks, not necessarily in your own organisation, but overall. It is always worth looking back at history. What the authors have done is to take data sets of risks that have actually occurred in the past in particular situations and they have then applied their anthropological technique to show how those risks are connected, how they have evolved from each other and how perhaps quite small changes in evolution might cause new risks to emerge. That is a fascinating technique.

This paper is at the cutting-edge. It deals with some of the issues of risk complexity. These techniques are no substitute for a basic ERM framework in the first place. The joint initiative that I referred to between the civil engineers and ourselves has produced a detailed comprehensive guide on how to introduce a full ERM framework. The ERM guide is on the Actuarial Profession's website (ERM – A guide to implementation, 2009). This deals with many of the key issues which arise when developing such a framework.

One of the important points that it emphasises is the need to look at underlying causes of risk. That is where the techniques within this paper can be used within such a framework, to dig down deeper into the underlying causes of risk for complex situations.

Mr J. E. Gill, F.F.A.: I have three observations on some of the practical realities and seek the authors' views on these.

First, in my experience of using risk appetites and asking people at the sharp end to model their processes, the engagement of those who are closest to the issues is absolutely vital. The difference in quality between situations where people feel the model was being imposed on them as opposed to where they own the model and are using it for their own purpose is significant. Therefore I seek the authors' views on this. With these complex systems, how do they make that happen?

Secondly, in any modelling process or mapping process, if you do not drive the mitigating actions coming out of the process then you might have a great answer but have you actually changed things? I seek the authors' views on when that has been successful using these techniques.

A third issue that I have come across, and Mr Gray referred to this in his remarks, is in terms of the engagement with boards and their tendency to use anchoring. For example, they like last year's risk appetite and therefore using a different technique and starting from scratch, which is likely to end up with a very different risk appetite, might be problematic. What experience do the authors have of being able to overcome the natural anchoring tendency that boards might have?

Professor N. Allan, co-author (University of Bristol): I will just touch on the point about the methodology and involving people in the model building. This process has to be viable. The concept mapping stage captures the language of the group which sometimes is a little different from the organisation. It is vital for them to really understand their own model. That is helped by keeping their language and not implanting your own language on top of it.

I would like to discuss two main points that Mr Lewin made. Firstly, just some clarification. In the paper we have used risks in the modelling using the evolutionary technique. They are risks; things that had not happened yet. We have not looked at losses in the paper although the technique can be used for both. Certainly, we have subsequently looked at losses and used the technique to anticipate future losses, such as rogue trading losses and safety related losses. Indeed some might argue that losses are a bit easier for people to relate to the concept of evolutionary emergence. The technique has been used to understand language, to look at identifying new strains of viruses and bacteria and for new product design in manufacturing; so it has a fairly wide set of usage now.

My second point is really an apology to Mr Lewin because if you interpreted from our paper that we are suggesting some sort of black box, we probably have communicated the approach poorly. In fact, one of the strengths of the Bayesian net modelling and the concept mapping approach is that it makes the model and assumptions very transparent. The fundamentals of the models can literally be built in front of boards or groups of executives, so they can participate and see how it is all put together. Bayesian Nets can get complicated when you start having many connections, but software these days overcomes many of the barriers of doing the maths. Moreover it is a theory that people already know in your industry.

Mr Cantle: Our experience of the Bayesian network has been that once you explain the background to how it works boards find it intuitive. The model retains the full transparency of why things are doing what they are doing, and boards can actually look at the model and see where those dynamics are coming from. In the context of operational risk, I think for a lot of boards it is the first time that they have had a discussion about operational risk in terms of business rather than statistical log normal distributions. The power of a Bayesian network is that it lays bare all of the dynamics that you are trying to communicate. Obviously, the skill that we have tried to articulate in the paper is how you need the cognitive part of the process to elicit what it is the experts want to get across and to simplify. Otherwise the experts will try to model the universe. Enabling the business to understand how they get from something very complex to a model which sufficiently captures the vital dynamics but which is understandable is a crucial step in what we have tried to convey in the research documents.

If you compare the approach to a classic actuarial modelling or a financial modelling exercise where you disappear into a room with a lab coat on and come back with a model, the board can understandably get grumpy that they did not understand where the model came from.

The cognitive process actually also leads to the question of how you drive the mitigating actions. By describing clearly the dynamics of the situation that they are considering the board are actually beginning to think in their own minds how they would deal with that consequence.

The anchoring point is something we all recognise. The nice thing about these models is that they are quite hard to fudge. Because they are quite complex, for the board to use last year plus 10% they have to fudge an awful lot of parameters in the model. What we found was that when presenting boards with last year's performance, this year's reality and enabling them to see what has changed it is much harder for them to revert back to the previous year's position. These are important communication tools. There is a lot of actuarial skill involved in the modelling. The art of the delivery is all about communication and the explanation of what is occurring.

Mr W. D. B. Anderson, F.I.A.: I should like to ask a question about the emerging risk part and the emergence of bubbles over time. You make a distinction between true black swan things, things that we do not know at all, and things that we do know about but are building up or accumulating over time.

If you cast your mind back to something like the dot-com boom of the late 1990s could these techniques have been used in businesses at that time in, for example, an asset management business? How do you think that behaviour would have altered, and would you have avoided the creation of some of these bubbles that we subsequently have seen?

Mr Cantle: One of the things that we try to be careful about is this. Obviously, we cannot claim that any of these tools would instantly influence someone's behaviour. If you tell someone that they are going to fall off a cliff and they decide to do it anyway, there is only so much that you can do. But what we have found is applying some of the tools to scenarios like you have described, with hindsight, there are actually signals there that you can see. Obviously, we have shown some tools in the paper, but there are other systems tools that you could use based on data, to find emerging trends underneath the surface that you may not have detected with your headline metrics.

The key thing is as a system goes towards a tipping point, like a bubble, it does give out distress signals which you can see. Whether they can be detected far enough before the tipping point to act is another question. Sometimes the answer is yes, sometimes it is no. Certainly in the current liquidity crisis you can see that in some of the bank data from 2005, three years before the crisis there was a lot of evidence of future problems.

Professor Allan: I will just add a footnote to that. We are advocating in this paper a number of tools that we have found useful to integrate together to tackle complex problems. One of the exciting things is that there is a whole raft of additional tools from complexity science and systems theory to help.

As Mr Cantle has just said, systems that tend to approach the point of criticality do exhibit certain behaviours. A bubble is an interesting question. I am not quite sure whether the evolutionary approach would predict a bubble, unless you had really good monitoring data. It also might not necessarily detect 'black swans'. It might tell you something interesting about how events are unfolding and would probably do so earlier than other techniques.

My main point though is on the concept of the 'black swan' – because I get a little bit upset with the 'black swan' label. I do not think a black swan represents a rare occurrence. The move from a white

swan to a black swan might seem extreme, but from an evolutionary point of view, we do have black birds. We have quite large black birds. So you might easily conceive it would be possible to have a black swan (indeed there are black swans), it is just that we might not have seen one. What would seem to me to be a rare coloured swan is an orange one with purple stripes and pink dots. That would be quite unusual because we do not see that anywhere in nature, yet! The idea of this technique is we would not spot the multi-coloured striped and spotted swan because we have never seen any of those characteristics. That is what I would call a very rare event that we would not pick up; whereas a black bird marking characteristic might evolve in a swan. Therefore I think Taleb should have called his book the purple dotted swan.

Mr S. J. Makin, F.F.A.: I was interested in your remark that your approach is not inconsistent with Solvency II. I wondered if you could expand on what you meant, commenting on the Solvency II applications of the approach. I am interested in particular in what needs to be done to make the framework consistent with the Solvency II Directive, in particular the calibration standards, and what the particular challenges are in doing this.

Mr Cantle: What we are getting at is when we have used Bayesian techniques in live cases you are really trying to get a handle on what your risks actually are. The danger if you come from a world where you are using just capital tools is that you tend to forget everything else. You do not have any obvious tools to use in understanding other risks. The kinds of techniques that we have shown here are just something supplementary that can be added to the risk framework which help boards think more deeply about the risk profile that they have and not just in capital terms. They also give a structured insight into the capital calculations that are needed. It is really around giving a broader perspective on the risk profile than obtained by just coming at it from a Solvency Capital Requirement (SCR) calculation aspect.

The Chair: You have talked about examples where you have used this with people. But is it actually being widely used? What is the reaction, if that is the case? How do people see this relative to the more simple approaches which they have tried to use in the past?

Mr Cantle: People are using those tools very widely, for reverse stress testing, scenario and stress testing, operational risk modelling and strategic risk profiling. The techniques in this paper are useful anywhere where you have a complex risk which you are trying to understand. If you are trying to understand how you should express a complex set of risks, whether operational or strategic or when you are modelling something really difficult Bayesian nets are a much better way of doing it, in our view, than just picking a lognormal set of parameters or trying to fit data that does not really mean anything because it is all historical and the business has changed.

We have seen companies over the last two years or so really begin to pick up on these things because as they get through the easy bits of Pillar I they start thinking about Pillar II for Solvency II. They have a lot of questions which their current tools cannot answer. We have found that a lot of people have seen some of these tools as a way of beginning to engage their boards and stakeholders around those areas. I am aware of maybe half a dozen firms who have actually picked up this paper and started using the approaches in it. I am aware of probably about another dozen that we have helped do things in those wider areas as well.

Interestingly, we have also helped mining firms, dairy firms and energy firms. The Actuarial Profession wanted some tools which we as actuaries could take more widely, we have done this, and the tools have worked very nicely in those wider contexts.

The Chair: At the end of the day what would success for this research work actually look like?

Mr Cantle: This work began seven or eight years ago when Professor Allan and I first started looking at systems in risk. It seemed to us it was a good time to make the profession aware of this so we responded to a call for research. Since then we have been doing a lot of other investigations into systems thinking that could be applied to risk that have not been supported by the profession. We have looked at applying systems thinking into areas like risk culture. We are aiming to understand how the risk framework might clash with the culture. A lot of people would say change the culture but that is not always possible. With an understanding of why the risk framework and culture clash you can, for example, seek to influence the framework to try and fit better with the culture. We have also done a lot of work around operational risk, stress testing and scenario testing. Boards seem quite interested in proving aspects of their models such as the correlations used. It is difficult to prove that correlations are correct. All of the data you have is around the mode. How do you prove that your tails are correct? Techniques like Bayesian nets can be used to try to evidence the logic of why things co-vary.

Professor Allan: The concept of complexity and systems theory is so well established that universities are realigning themselves to try to provide this to undergraduates and postgraduates. The University of Bristol Bath system centre, which is funded by the Engineering and Physical Sciences Research Council (EPSRC), has 70 engineering organisations with doctorate students. This shows that the work has gone beyond just the education of students. Systems behaviour is something that all sorts of industries are thirsty to understand and eager to start to apply as organisations realise that the reductionist paradigm in much of our thinking in engineering and finance, needs to be supplemented.

There is a continuing research into the philosophical debate around some of these topics. Certainly the debate about emergence itself which started in 1930 is still very active in many different disciplines.

This is quite an exciting time as complexity science, biology and systems theory combine. We have not seen the half of the techniques, tools and understanding that are yet to come. We have just scratched the surface. Even the tools we have proposed in the paper have already been improved and developed considerably since the time of writing it.

The Chair: What would you consider to be ultimate success for this strand of research?

Professor Allan: Success would be that practitioners in the Actuarial Profession are confident in using complex system tools alongside their existing traditional tools.

The Chair: We have heard from Mr Cantle that there are a number of companies that are in that position. I guess there must still be a lot of companies that are not. We heard at the beginning from Mr Lewin that in some cases it is still a question of putting in place a robust ERM framework before getting into some of the more complex areas here.

Mrs I. M. Paterson, F.F.A.: Could you share with us any examples of the work that you have done that illustrated a surprising effect using your deeper analysis?

Mr Cantle: We find that the whole point of the cognitive mapping process is you are holding a mirror to a bunch of experts. Whatever they discover, they already know, but they possibly have not

articulated it in a consistent way to each other. If you have organisations which are very slick at communicating, and they speak the same language, often they would not have surprises because they already do it.

A lot of companies are busy. They do not always use the same language in a technical team as they do in business. Wherever there is a boundary on a discussion, often you will find that those experts will learn something that they had not seen before because of the interaction across the boundary.

We had an example in one of our very early engagements on this about six years ago where there was an organisation which is a US mutual. They were telling us lots of interesting things about how their strategy would work and where success comes from. We asked them the unthinkable question about what would happen if they were no longer mutual. Around that time there had been examples of companies that had had to go down that route. But they just could not discuss it because it was a belief, a value, which they did not want to talk about. Actually, as it happens, we were doing this over a video call. They started realising that a lot of their strategy would unravel very fast at exactly the point where a capital call, for example, would cause them to demutualise. So there would have been a massive cascade of catastrophic strategic failure from every single angle. There are things like that where people will not discuss something until they see it laid bare through a cognitive mapping process and then they have to start actually acknowledging that there is an elephant in the room and start talking about it.

We have had other examples in reverse stress scenarios where people find out something very simple, such as that annoying the shareholder could unravel the strategic viability of the business. That is very pertinent when your shareholder is a family. There are simple things where what you say could upset the family. You say report a string of bad news, nothing big, and because they are just a family they get upset and they try and divest the business. These are the sorts of simple things which are perhaps unsaid on a day-to-day basis but once they are there, written down in a cognitive format, cannot really be avoided.

Professor Allan: I have had similar experiences. Sometimes they are quite embarrassingly simple. The one that jumps to mind is a company and their reward system for their agents selling insurance policies. The strategy of the company was producing innovative products. What they had not realised is that the behaviour of the insurance agents in the field was that they would do this particular company's business up to a limit and then move resources to another insurance company, because after they had reached their target their commission was dropped substantially. So the agent's profit maximising behaviour was to go to some other company. This was exacerbated by the fact that the insurance company's core product was very attractive to customers and so the agent's target was met easily with the core product and there was no need to push the innovative new ones. So their innovative products, even though they were very innovative and designed with a lot of stakeholder engagement, were simply not selling because at some point their agents stopped selling them. Why would you push a new product when you can sell your existing one so much easier and get the quick 5% before going on to 2½%? In the Board room that knowledge was there but nobody had actually connected it. It became a huge revelation to the whole group that this was actually their blocker in their innovative products being sold. Unbelievable maybe but actually this loss of connectivity is not uncommon. That is the sort of occasion you remember and say to yourself, "Yes, this has made a difference." Ever since then, that company has really tried to understand what is connected and what the feedback loops are. If there are two or three interacting feedback loops it is very, very difficult to understand its behaviour. So articulating it, putting it on paper and sharing it, is actually a very good way to try to understand the dynamics of the situation.

Success then would be if people understood the feedback and the dynamics of what they are involved in and how it impacts their behaviours.

Mr J. Fraser (Guest, Head of Development at Registers of Scotland): I am a little bit perturbed by some of the references to intuitive understanding. Mr Gray was mentioning that it would be better if directors had some kind of intuitive explanation of what was going on. Mr Lewin also referred to the paper exercise as being the one that works and the Bayesian stuff can back it up. My understanding of this whole world was that by doing the modelling you would actually bring out unintuitive behaviours, the real behaviours. I thought the power of it was that it would teach us that intuitive thinking was perhaps not sensible and you would learn from these models. Do you want to comment on that?

Mr Cantle: The purpose of trying to introduce these techniques is that they do challenge you at a fundamental level to start thinking beyond what you intuitively know. The cognitive exercise examples we have just been giving are where the intuition has got people so far. Then confronted, effectively, with a cognitive model, they have been able to push each other further, which none of them would have done on their own.

Professor Allan: I would probably say the example I just gave was non-intuitive.

Mr I. Campbell (Guest, Chief Information Officer at Registers of Scotland): When I viewed Mr Cantle's Bayesian system slide with all its percentages and bar graphs, I had a visceral reaction. The hairs stood up on my arms because I can just imagine some people going, "Oh, my goodness!" and being utterly beguiled by it. That to me is a serious risk in itself.

Board members are not perfect. I struggle with my fellow board members. I am not an actuary so I am less numerate than many of the audience. Computing science was my background so I am significantly more numerate than my fellow board members. I struggle, however, to discuss models of risk and the interconnectedness of it. With my background of research in computer science, I built some Bayesian models and used Bayesian techniques, and I saw nasty, complex almost chaotic behaviour in them with anything apart from the most simple models. My concern is that building on those models gets more and more complex. The model itself becomes a complex system that is difficult to predict. Therefore I would not trust it and I would not want others to rely upon it. But it is so utterly beguiling because of the apparent detail and control.

So my question is: when do you know when to stop with anything but the simplest of Bayesian models? My feeling is do the techniques, build the models as an exercise and then throw them away and use the lessons learnt.

The Chair: I should like to give the authors an opportunity to answer that question and to give any closing remarks which they have.

Mr Cantle: The techniques we have tried to show are meant to facilitate the thinking process rather than creating the classic actuarial black box although occasionally we may have been guilty of that. Essentially, they are a device for calibrating and understanding in the risk department how factors interact and communicate with the business. The key thing is taking the insights that the techniques give you and then presenting that back in an understandable way to the board. There are two different things here. There is a power point slide that goes to the board with the explanation

and then there are the devices that we have tried to show about how you actually calculate those things. They are not necessarily the same thing. A model is always wrong, as they say. Sometimes they are useful. I hope these models are useful.

Professor Allan: By wrestling with the initial complexity, you get a degree of simplicity that emerges because of the patterns that emerge. There is a scary point when you are doing some of this modelling when you think that it is not telling you anything. You need to have a little bit of confidence that the patterns will come through, it is just that you may not be viewing it at the right level of detail.

On concept mapping I absolutely agree. A full concept map with maybe 200 nodes does look like a dog's breakfast. You would not show that to a board, even though it may have been elicited with the board's input. There is an analytical basis behind even a concept mapping tool that will allow you to look for the key nodes. You can quite quickly collapse down on those nodes without losing the data behind them so that the key story, in systems language the key modes of behaviour of that system, starts to emerge. That is one of the key issues for practitioners.

The model that you build after the elicitation with the board or executives obviously gets fleshed out with different levels and there are often lots of repeats. You do not need to show that. The story, the narrative, is kept at the high level. That is not a sleight of hand. The evidence is still there, but you are making it presentable.

My closing remark, is to thank our other authors who could not be here today, Prof Patrick Godfrey and Dr Yun Yin. I would on their behalf like to thank the Actuarial Profession for allowing us to do this work, to take the courage to embark on something new like this. I should like to thank you all for the opportunity for myself and Bristol University to be involved in this important piece of research for the profession and other disciplines.

Mr C. I. Black, F.F.A. (closing): I would like to thank the authors for a very interesting paper. It challenges us as a profession to add to our toolkit for addressing risk appetite and emerging risks. The introduction through the examples helped prepare the reader for the sort of thought processes that would be required in the more complex examples later on the paper. This encapsulated one of the messages in the paper that we need to find new ways to think about the complex issues in this field.

Some of the examples in section 3 on emerging risks were helpful reminders of the timescale on which change can emerge. The example of the emergence of personal computers reminded me that the timescales on which change emerges can be long. This is particularly the case when so much of our thinking is devoted to ensuring we have models that meet the required standards for 99.5th percentile outcome over a 12 month horizon for the purposes of setting capital. Recognising the time when a risk moves from being a distant prospect to a clear and present danger is always challenging. Consider the example of the Hillsborough revelations last week, these make clear that some people must have known what became apparent more widely last week. This helps us understand that the quality of our understanding of a risk is very dependent on the information we have available.

On risk appetite as I read the paper at times I felt that the approach was that this was something to be discovered. This jarred a little with my thinking that risk appetite is something that should be

decided. However, as section 6 highlights the tools are used first to develop understanding of the risks and then to aid in the process of deciding how much risk the firm wants to run.

This is not something that is easy to do and tools that help make this link are always welcome, and therefore worthy of greater consideration by actuaries.

The prize of an understanding of risk evolution offered in section 7 is one worth having. One of the messages is to let the data guide you to the answer although there is still a lot of work to do to get to the data and then help it guide you to the answer. The key themes that arose for me in the discussion are:

Model complexity- we heard that these models are hard to fudge. Mr Fraser noted the challenge of existing modelling and Mr Campbell warned of the risk of relying on a complex model.

- The role of boards of directors- the challenge of getting their buy-in and overcoming what Mr Gill referred to as anchoring in last year's risk appetite answer.
- Cognitive maps and Bayesian networks- the cognitive step maintains the link to the organisation and picks up interconnectedness. The ability to expose the elephant in the room was highlighted. One of the clear messages that we have to take away is the use of the model as a communication tool. The warning that we should not show it to the board, at least not without a proper cover on it.

There were also some areas where there might be discussion but there was little or none:

- I thought the role of experts and expert judgement would feature more in the discussion.
- Behaviour and the potential for people to 'game' the model once they understood how it worked was not challenged.
- Very little was said about phylogenetics so perhaps this is an area that we have more to learn as a profession.

The references to Darwin in the paper reminded me of a quote. It is a quote of a quote because I first used it after reading it in the Berkshire Hathaway Chair's report, so it is a bit like some of those risks from the phylogenetic analysis. The Darwin quote is "Ignorance more frequently begets confidence than does knowledge".

I thought this was apt especially as it was made as a warning about derivatives disclosure after the 2003 losses by Freddie Mac, my recollection is that this was an unexpected risk at the time but is it one that could and should have been discovered earlier?

Wikipedia tells me that what Darwin is referring to is known as the Dunning-Kruger effect after they did work on this when they were both at Cornell University:

"a cognitive bias in which unskilled individuals suffer from illusory superiority, mistakenly rating their ability much higher than average. Actual competence may weaken self-confidence, as competent individuals may falsely assume that others have an equivalent understanding."

As Kruger and Dunning conclude, "the miscalibration of the incompetent stems from an error about the self, whereas the miscalibration of the highly competent stems from an error about others".

The authors have contributed work that encourages us to shift along the spectrum from ignorance towards knowledge. As a profession, we should reflect that the combination of actual competence, even at the price of lower levels of self-confidence, is a direction we should be willing to travel.

The Chair: It remains for me to express my own thanks and the thanks of all of us to the authors, the opener, the closer and all of those who have participated in this discussion.