

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

Abstract of the London discussion

[Institute and Faculty of Actuaries, 28 November 2011]

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 R. C. Dix F.I.A.): Welcome to you all to tonight's discussion on A Review of the Use of Complex Systems Applied to Risk Appetite and Emerging Risks in ERM Practice. I should like to introduce Neil Cante to introduce the paper.

Mr N. J. Cante, M.A., F.I.A. (introducing the paper): The two challenges we were set by the profession, are to work out how we should define and use risk appetite and then how we should help people identify and assess difficult risks. Risk appetite is about how you control your business in order to maintain the uncertainty of your business outcomes within some boundary. A good first question is: how do you know what the outputs and inputs are doing in terms of how they are linked? And then: How do you quantify the current level of risk?

Within our industry, a good answer to that from one angle is the capital base. That is not universally true across all industries.

For emerging risk, the quote of Donald Rumsfeld springs to mind, about the “unknown unknowns”. That is usually a challenge with emerging risk. Taleb introduced the concept of black swans, so now we are all hunting those everywhere. It comes down to a question of “How can I see what I do not know I should look for?”. That was the way that we started thinking about emerging risk.

In looking at risk, certainly at the enterprise level, a few things come to our attention. First of all, holistic views are everything. You have to look at the whole to understand what is going on underneath.

What you are looking at has a knack of changing as soon as you have looked at it. So adaptation is pretty important. Everything is very non-linear. So things have feedback loops and are very counterintuitive. Everything has people in it which means the cognitive biases are pretty much everywhere.

The study of systems gives us a way of understanding and thinking about such complex problems. Companies here are characterised as complex adaptive systems. Here we are looking at open systems because we are interested in the movement across the boundaries between the company and its outside environment. Such systems have been studied for a long time. The good news is that they

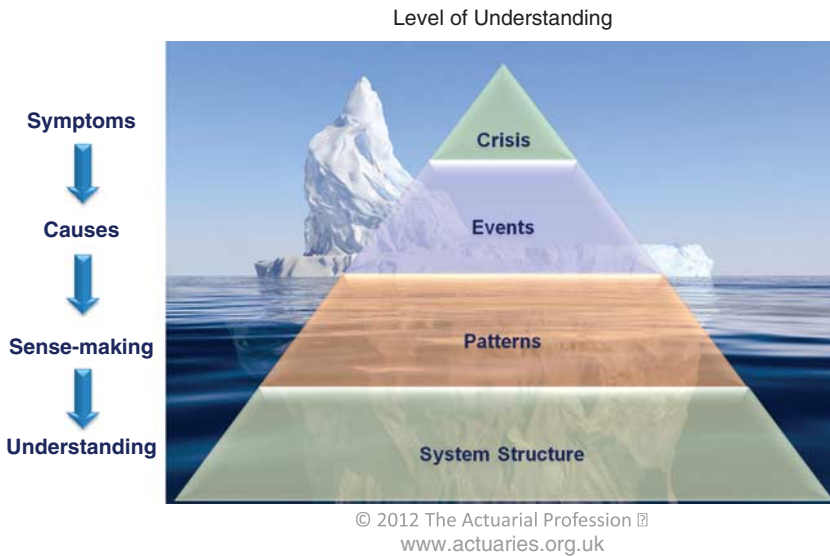


Figure 1. Level of Understanding

have some key properties which we can leverage in our understanding of risk. Three in particular which are relevant are: “Emergence” meaning that what comes out of the system cannot be known by simply studying what goes in; “Adaptation” meaning that things change on a regular basis; and “Nonlinearity” meaning systems have features like feedback and feed forward.

Systems have been studied widely. They are not restricted to the actuarial world. Essentially, this research paper is about bringing some of those into our context so it is easier for us to understand how to apply them.

The system tools we are looking at are going to help identify and understand emerging properties of the systems. Particularly for the emerging risk part of our research that is very important. We describe how the system works in terms of all the key underlying interacting factors, particularly in the risk appetite part where we need to understand how the inputs produce the outputs. We link together lots of different perspectives. We bring together what people know with what they have measured and what they have said.

This picture in figure 1 explains the difference between traditional risk efforts and a systems risk effort. If we start from above the water line, we can see a crisis and we carry out some kind of causal analysis to try to work out the events that led to that crisis. Quite often, risk systems stop there and we catalogue things, count things and model them with distributions. But we still do not know why it happened which is why the emerging risk problem is still challenging, because we do not have the insight about the underlying mechanism. We need to take a systems approach so we can see below the waterline and start spotting patterns. Ultimately, we get to an understanding of the real mechanism behind what is going on. That is what we are trying to achieve by taking a systems approach.

How do we apply this approach to risk appetite? We worked out that there is a distinct lack of consistency in the use of certain jargon, so we started by defining what we were going to mean by our terminology so we have a benchmark.

The first definition is uncertainty, which is a lack of certainty. It means that you do not know exactly what is going to happen. Risk is where one of those uncertainties is negative from somebody's perspective. We do not subscribe to the phrase "upside risk". As far as we are concerned, risk is always negative and the upside is that you get more than adequately rewarded over some period for taking the risk. We have tried to be careful in how we have defined risk appetite. We have not bought into a specific industry view. We have tried to keep it broadly based around the uncertainty in delivering the objectives that a firm has. This permits actuaries to use these sorts of tools in other industries, not just in insurance.

By risk limits we mean how you constrain the operational activity so that you keep the outputs within the risk appetite that has been set.

Risk appetite starts at the strategic level by agreeing what it is that the company is trying to achieve. So there are a few key goals. Usually there should be some planned outcome for each of those goals, so the board would sign up to that plan. However, life is not very predictable so the board needs to explain what outcomes they would tolerate. That specifies a floor to business performance. More than that, the board needs to say how often it would be acceptable to fall to that level. So we have done three things. We have now expressed the uncertainty that we are prepared to accept around achieving our business goals. More than that, we have to drill down for a risk appetite framework and say, "where do those sources of uncertainty come from and which business drivers lead to those sources of uncertainty?" Through an understanding of how all of those things interact we should be in a position to try and control our organisation to produce outcomes with a particular degree of uncertainty at the top. This understanding has become a tool for knowing how much risk capacity we have and the allocation of our resources to try to deliver the business goals of the organisation.

From a systems perspective we have lots of things going which we are trying to constrain by putting limits on them. They are going to interact through the processes and operations of the organisation and produce outcomes which we want to keep within certain boundaries. This is a large, complex multi-objective optimisation and control challenge.

In looking at risk appetite we have tried to find a solution which copes with nonlinearities. We looked for an approach that can adapt and learn, because that is necessary given the reality of the situation it will be used in. It needs to be communicated effectively because if no one understands it, it is pointless. It needs to suit a wide range of firms so that we, as risk professionals, can use it in many situations.

We identified an approach, which appears to link the business drivers to the uncertainties at the top. In our research paper we have proposed a combination of cognitive and data-driven methods which help us leverage both what the experts know and what they can measure. The technique also fits with forthcoming requirements like the Own Risk and Solvency Assessment (ORSA). For insurance firms clearly that is important. It also fits with broader planning requirements and strategic thinking. The model remains in the language of those who contributed to the building of it. That is useful for the embedding and use of such a model. We have checked that it does not conflict with anything in Solvency II.

As a first step we use a cognitive map to capture what experts know about where the uncertainties in business performance come from. By distilling that and combining all of those multiple perspectives into a cognitive map, we are able to reduce, or eliminate, the biases in how they would



Figure 2. Sources of risk

normally express that. It is also a nicely structured format for capturing what they know, which yields nicely to mathematical analysis so that we can start simplifying it without losing any of the critical complexity in what they were trying to express. That is the key point. It is simpler but it is minimally complex. It still retains all of that nonlinearity that the experts were trying to explain.

Once done, we can translate that into a Bayesian network.

Figure 2 shows at the top three business goals. In the case study they were profit and loss, balance sheet and reputation. We then go through two layers of sources of uncertainty, in this case study organisation using some of the Solvency II sources like market, credit, insurance and operational, and so on. At the bottom we have a series of business indicators which inform us about the levels of those uncertainties. One of the nice properties of Bayesian networks is that they can propagate evidence up and down. If we establish, for example, our level for risk appetite at the top and input that to the model, that propagates down to tell us what the limits underneath should be. We have an example of a credit risk setting exercise where we have input our desired output at the top and the model has shaped itself to tell us what the limit should be on things like credit rating dispositions, exposure to counterparties, and so on.

You can also propagate evidence up, so it becomes a monitoring tool as well as a setting tool. You enter observations in the bottom part of the model and let that propagate up. It gives you a guess as to what your current level might be. That actually enables you to ask questions of your other risk tools to see whether they agree.

In summary, the proposed approach for risk appetite does embrace systems thinking. We have started holistically looking at the problem before breaking it down. It is very scalable. It works on small, simple things and very big, complicated things. It works on any type of firm in any type of industry. It reacts to new information, so it adapts and learns as you find out more things about the sources of uncertainty in your business. It provides a basis for both setting and monitoring within the same tool. It uses expert knowledge until data is there to support that, or otherwise. It enables you to make progress in areas which are very difficult and as yet not necessarily monitored. It is

used in a way which remains in the language of business people, so it does not get a black box format. Essentially, it translates risk into business terms. It actually helps you embed the risk appetite within the organisation.

For emerging risk we looked at biology, a different science. We thought that it would be interesting and instructive to think about risk as an evolutionary process. In the paper we have documented why that is a reasonable and sensible thing to do. It gives us a classification system so that we can understand our risk profile better. It tells us where all of the dynamics might be and where they might go in the future. It gives us a different risk profile rather than the risk capital profile that we are used to within our industry. Looking into other industries gives an easier language to understand. Cladistics and phylogenetics, which are studied in biology, are useful for classifying things with many characteristics.

We looked at groups of things which are similar by virtue of their characteristics, and, once we have found them, we can see their ancestry and their evolution. We looked at the similarities and differences between those different groups and once we have understood how they have evolved, we can ask some interesting questions. Where is evolution most prolific? That gives us some clues around emerging risks. How did they get there? Which things evolved together? Which characteristics are most active? That might give us clues as to things again which might emerge in the future. It gives us a basis for creating scenarios for emerging risks in a very structured way.

One feature here is that we wanted to build a technique which leveraged something most firms would have.

Most firms, both in our industry and outside, would have some form of risk register where a number of risks that might exist in the organisation are listed. Typically people classify them using one label which is, arguably, an oversimplification. To apply this technique you simply use more than one label to classify each risk.

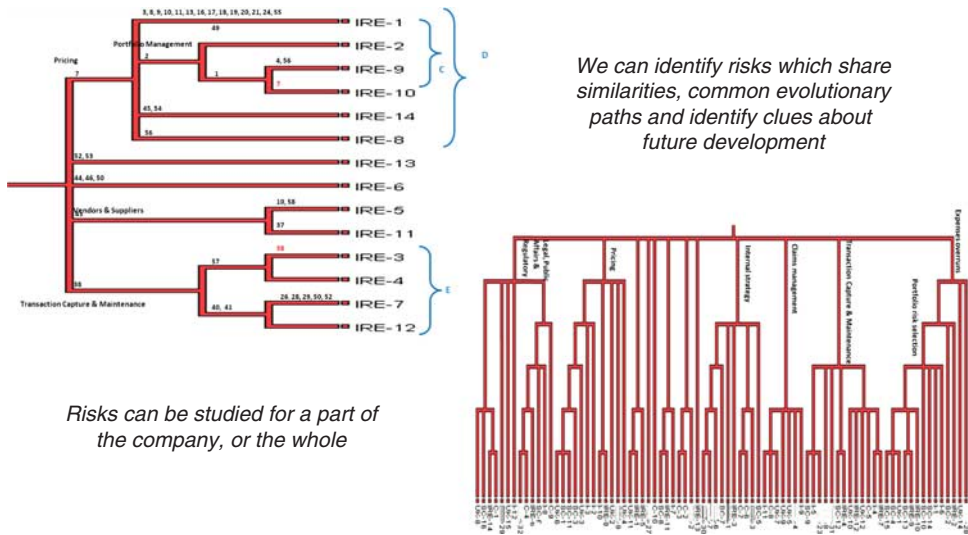
Once you have done that, you draw the sorts of things which you can see on figure 3. In the case study in the paper we used a multinational group as the basis for the case study. We were able to plot out individual businesses as well as the group overall.

On the left of the figure you can see one of their business units very clearly has two big clusters of risks, which seem to share some ancestry in terms of their characteristics. That is around pricing and servicing.

Looking at the bottom right, for the group overall, there are about half a dozen things which come to light when we look across all of their businesses. Visually you can pick up those patterns. This approach helps us spot patterns in the data, which we might not otherwise see. This leads to asking some questions. If we look at the shape and we see lots of bifurcation, this tells us that things are changing a lot and so it could be an area for emerging risk in the future. Which branches have characters which are changing a lot? Which have the most characteristics? These are things again which might adapt in the future. Which characters themselves evolve more frequently? That might give us a clue that if we see such a character, we should be thinking in terms of that changing in the next period.

Overall, we might spot patterns about characteristics that show up together or in sequence. That gives us clues about what we might see next once we have seen one of those happen in the past.

Evolutionary Risk Profile



Risks can be studied for a part of the company, or the whole

© 2012 The Actuarial Profession
www.actuaries.org.uk

Figure 3. Evolutionary Risk Profile

Again, we have structured it, which means that we can use mathematical analysis to try to find out things that we cannot see visually. There are lots of analytics on those types of trees to try to find more insight about how things are structured.

In summary, we find that modern risk is a complex beast and therefore needs a holistic systems view to try to make sense of it. That is exactly what systems science is all about. The techniques that we have proposed in this paper are things that leverage existing infrastructure, so they do not demand a wholesale build of new things. The tools can indeed, and have been since the paper was written, be applied to any type of firm in any type of industry.

The techniques make the areas of complex risk transparent. The authors do not claim this will give you great insights into true chaotic black swans. It might, however, help you spot those things as they become ordered again. These tools are designed for complex risks not chaotic ones. These techniques also leverage a lot of the skills of our profession. In terms of understanding business dynamics, and in translating those into modelling tools, they very neatly fit into the skills that we have developed.

Mr A. D. Smith, student (opening the discussion): This paper comes at a slightly difficult time for anybody involved in prediction. We are in the middle of a financial crisis that not many people saw coming, so a paper that gives us some new ideas for modelling emerging risks is welcome. It is also very helpful to have a broader perspective as to how other professions outside the Actuarial Profession treat these problems, and how they discuss them.

I did some research in preparation for this. There is an organisation called the Systems Dynamics Society. They have been organising annual conferences on this since before I was born. One of the

hot topics in the late 1970s and early 1980s was using complex systems to forecast lapse rates on variable annuity products. It seems that these remain relevant today and probably are still not solved.

We are all used to using statistical techniques, and maybe we are most used to using them on outputs of complex systems. This paper challenges us to think about whether we have really captured what is going on inside those systems. For example, we might try to model losses that have come from operational failures by looking at a history of past losses and fitting some sort of process to those. The danger is that by not going inside the system, and understanding what is generating those, we risk constructing some sort of model that prejudices the outcome. You might say: "This looks like some Poisson process and some Pareto losses because I have read this in a mathematical paper." The paper tonight offers us some better ways of addressing that sort of question.

I endorse what the paper is saying about modelling the components of a system taking data, not only on the losses that emerge, but also on what is going on inside and what might ultimately lead to those.

The question that you might want to discuss is whether you think that this is really new. I guess many of you here will have written computer programs to project balance sheets and profit and loss statements, and you might have built them up from lots of things, from policies being written, from policies going off the books, from mortality, from management behaviour, from investment returns etc. Many of those principles are fairly similar. What we have learnt is a new language and perhaps some new ideas, but not something that seems to be a radical turnaround on what we had before.

This paper contains many tools. The section on risk taxonomy is especially interesting, as is the information about cladistics. People have been using cluster analysis in actuarial problems for quite a while, particularly for general insurance. The approach that struck me as interesting was the way in which the data organised the way you think about it. So, rather than starting off with a predetermined classification of kinds of losses or kinds of events, you start with your risk register and apply some of these classification techniques and let the data tell you. So the theory is that this would help you to spot emerging trends more quickly.

There are a lot of diagrams in this paper: some I found useful and interesting; some of them a bit abstract; some, frankly, baffling. Is the way I consider useful versus baffling figures an indication of the way I elicit judgements from experts and if others differ the way they elicit judgements. The danger is that when we go to experts to get judgements, we try to make the experts conform to our model rather than the other way round. I will give you an example: you are trying to run a stochastic model of risks affecting an insurance company. In order to make the model run, there is a box which says "Insert copula correlation matrix here". You have a great big array of numbers to fill in and you are not quite sure where you are going to get them from, because there does not seem to be any data. You say, "Aha! I will ask an expert". So you find an expert whose real area of expertise is natural catastrophes or marketing. You go to him or her and you say, "I have a box to complete in this form and my code will not run until I have filled it out. So, can you use your judgement, please, to tell me what number I will put in the box?" It seems to me that unless your expert is also an expert on copulas, there is not a lot of chance of getting a sensible answer to that question. It is like trying to get intelligence under torture. Maybe some of the diagrams that we have seen, although they might not be helpful to you, could be helpful to different sorts of experts when you are trying to get information out of them.

These applications seem to require a huge amount of experience and knowledge as well as mathematical techniques. How useful the model is will depend greatly on how well the person who built it understands the system that they are trying to model. When you hire somebody to build a system like this you are reliant on their judgement and expertise, much more reliant on that than you are on any particular software tool or procedure. The problem is it is quite difficult to work out the extent to which somebody understands the system if you do not fully understand it yourself. If you are trying to find somebody more expert than you, how can you tell if the expert knows anything? I wonder whether some of the cognitive mapping tools that the author described earlier would be helpful in helping us to understand how much somebody understands. Maybe some of those diagrams that look like spiders' webs could help with that. Also when it comes to certifying experts, this profession does quite a lot of certification. We have put a lot of effort into certifying skill and experience, and things like our CA2 modelling course and our Chartered Enterprise Risk Analyst (CERA) qualification are examples of that. Maybe the Actuarial Profession will in future be certifying complexity systems experts.

Mr J. P. Ryan, F.I.A.: The traditional risk management approach that the authors described at the beginning of the paper talks about limiting risk of loss. In fact, speaking as a member of the Institute of Risk Management, that is not the case. It is identifying and mitigating and controlling the risk. You may not limit it but you will do your best to mitigate it and identify it. A good example of that is sprinklers do not limit the risk. A building can still burn down. But they clearly mitigate it. Even with perfectly working firewalls there are a couple of cases where the fire has blown material out of the window and the fire has gone round the firewall. So it also does not limit it.

The authors state that computational speed is one of the important issues. This is not necessarily the case. Indeed, the effort on computation speed is one of the reasons we had the banking crisis in 2007. Because of the need to compute VaR risks quickly as trading floor risks change daily, computational speed was very important. In order to do the calculations sufficiently quickly people made approximations, which worked quite well on a daily trading basis. When integrated into the full model, it just missed the whole point. Everything had to be done quickly, including the whole model, and they missed the fact that credit risk and other risks were much more compact and had curtailments. You may need to look at computational speed but a lot of this type of analysis is going to take a long time to do. A lot of it is not going to change very much from year to year, so in many cases accuracy is more important than the other.

The phylogenetic point is a very interesting one. There is an extra way you can look at emerging risks. One of the problems that Darwin had with his own theory of evolution was that it did not explain the existence of music and why musical ability appears to be a characteristic that was emerging but did not. This was something that Darwin flagged up at the time.

I recently raised this issue with a well-known behavioural economist. He was talking about the importance of lifestyle economics, which is a big growth area at the moment, and it turns out that musical ability is actually closely related to our ability to co-operate with each other, and co-operation is one of the things that has led to the advance of human beings. This therefore provides an explanation as to why evolution can explain the existence of musical ability.

Applying that logic to the paper means you might want to map how the characteristics across an organisation are actually changing, because that might give you some clue to what some of the emerging risks are. An example of that in practice might be employee satisfaction. Not an obvious

one, other than possibly staff turnover, which is probably not a major risk. Actually, something happening to employee dissatisfaction can sometimes give rise to fraud losses, and other things going on. So it is worth putting an extra step in that process looking at characteristics across the organisation and how they might change. That might give you some clue to what new risks come through that we have not thought of – Rumsfeld’s ‘unknown unknowns’.

Turning to boards of directors, I agree that it is important to get their input, because the non-executives in particular need to consider what the model is actually missing. How rigorous is it? This was one of the characteristics of the banking crisis. None of the non-executives picked up half the issues that they should have done had they been doing their job properly. A lot of this needs to be translated into terms that they can understand and ask intelligent questions as to whether all these things have been picked up in the causal models?

Dr L. M. Pryor, F.I.A.: I welcome this paper as it is very useful in providing actual practical tools for people, tools that we can go out and think about actually using.

I would particularly commend the use of Bayesian networks. One of the things to remember about risk appetite is that sometimes you want risk. Risk is not always to be avoided. There is definitely a positive side to it, too. Using Bayesian networks can help you spot the ups as well as the downs.

I have a big question on phylogenetics (which is absolutely fascinating). I’m afraid I do not have an answer to the question, but would be very interested in hearing other people’s reactions. Phylogenetics works in biology because species are, in a sense, real things. There is a definition of a species from outside phylogenetics, in terms of individuals that cannot successfully interbreed, or whatever. Although there are several different definitions, people tend to be reasonably clear on what constitutes a species. Most of us know that evolution can really happen at the species level but the fact that new species can evolve does not necessarily imply old ones dying off. But do risks really evolve in the same way as species do? And are risks in that sense real things? Are they capable of being defined from outside this technique? In many ways I think not. I think that we can create a new risk or a new division of risk simply by naming it, simply by bringing together a collection of characteristics and saying, “Okay we are going to call this risk Pat or something”. In other words, just by saying, “This combination of characteristics is important therefore it is a risk.”

If you had talked to any of the Italian merchants of 500 or 600 years ago they would have understood the concept of liquidity risk once you had explained it to them, but would not necessarily have given it that name. Indeed they might well not have had a specific name for it. So, how much does the phylogenetic technique depend on the evolutionary analogy or is it just a very, very useful classification tool for existing combinations of things that we have chosen to call risks?

It would be extremely useful to spot links and relationships between risks and possibly spotting that risks that we have thought were diverse, were in fact quite closely related. I suppose that two risks that hitherto had been considered quite separate could in that way be thought of as a new risk and be given a new name. But is that really evolution? That is one question.

The other question is: how much does the technique depend on the evolutionary analogy or can one throw away the evolutionary analogy and just use it as a valid technique in its own sense? I would be interested in hearing other people’s thoughts on this.

Mr C. G. Lewin, F.I.A.: I have had the privilege of leading the joint initiative between the Actuarial Profession and the Civil Engineering Profession on risk, and particularly the management of risk, starting with projects and then moving to strategic risk and other kinds of risk, including now enterprise risk. Two of the authors of this paper, Mr Cante and Professor Godfrey, have also contributed to the work of this initiative. So I was particularly interested to see what came out of this paper.

The model I have in my mind is very slightly different, subtly different, from the ones which apparently the authors have. I always envisage two systems, one being the world, and that is a changing world. There is a whole lot of information that you can get about how the world's situation has changed, how various factors have evolved, how they will evolve in future. That is not of course complete information. There are many gaps but there is lots of information out there which one can acquire.

In my model there is also a second system, which is the company, or fund, the entity, the enterprise, which you are considering. That is a system in its own right, consisting very largely of people but also other resources, as well as relationships with suppliers, customers, and so on.

What you are really trying to understand is the system which comprises your enterprise. I emphasise understanding as deep understanding is really important. Some of the techniques set out in the paper might help in that understanding. I stress "help", because the experience of the people who have been concerned with the enterprise for a long time will also be a material factor in helping to understand the enterprise's system.

If you think about these two systems, and then you think about where they interact with each other, not just now but in the future, that is where there will be all sorts of pressure points on the enterprise, and also opportunities facing the enterprise from time to time. Understanding these points of interaction is something which needs to be fully explored in addition to the points referred to in the paper. In particular one should seek to understand where these points of interaction are, who are the people in the company who are really important in controlling those interactions, and what are their attitudes to those interactions.

What one is concerned with in thinking about enterprise risk management is the future variability of the results from the enterprise and, as Dr Pryor said, that means the upside as well as the downside. In fact, some outcomes, which you now think of as likely to be unfavourable, might actually turn out to be favourable, and vice versa. So it is important to study variability as a whole.

I thought I detected from something said in the introductory remarks that the paper was only looking at downside risk. However, variability as a whole is really important. By taking that philosophical approach, one will also then be looking at the opportunities generated by the interaction between these two systems I have described and trying to maximise the extent to which the enterprise might be able to prepare itself to benefit from those opportunities.

A technique one can apply to improve understanding is to have a methodical search for additional knowledge to reduce your uncertainty. That means first of all pooling the knowledge of the people in the company because "A" will have knowledge about certain things and "B" will have knowledge about others, and so on. It also means a methodical search in many other ways, using the Internet, using research papers, etc.

And finally on the question of evolution of risks, which Dr Pryor referred to, I think that some risks do evolve. I do not think that there is any doubt about that. The history of the enterprise and perhaps the kind of techniques set out in the paper, might help to understand this. But I certainly do not think that *all* risks may evolve. Some risks are totally new. They are on the horizon perhaps. You might or might not be able to see them. But they do not really evolve from any direct developments within the enterprise. You might be able to see them evolve in the world as a whole. or see them affecting your competitors. You might be able to predict perhaps that they will have an effect on the enterprise. But I am not sure that the kind of techniques set out in the paper would really help you in that particular area.

To conclude, I am not entirely sure about the model set out in the paper, whether it is sufficiently comprehensive. The techniques set out could in some circumstances be helpful in improving that understanding, but they are only part of an overall process of risk management.

Mr Michael Thompson (Guest, International Institute for Applied Systems Analysis, Austria): Mr Cantele, I noticed very early on you mentioned “certainty, risk and uncertainty”. That goes back to Frank Knight in the 1920s. He said if you know for sure what is going to happen, that is certainty. If you do not know for sure but you know the odds, that is risk. And if you do not even know the odds, that is uncertainty.

That scheme is fine as far as it goes. The trouble is it hardly goes anywhere, because most of the problems with risk that we are dealing with do not just involve uncertainty in that Knightian sense, but they involve contradictory certainties, that is completely mutually contradictory convictions about how the world is, what is out there, what the risks are in the world.

For instance, I did some work on siting liquefied natural gas terminals, and the big question as to whether a site is safe or not is how far could a spill of liquefied natural gas on the ocean, under ideal conditions, spread and still go up in flames? Some experts said half a kilometre is the maximum. Other experts said 50 km. If the first expert is right it, you can put the terminal almost anywhere, for example, Rotterdam harbour was favoured at that time. If the other expert is right, you can hardly put them anywhere.

Another example would be how much of Europe’s energy requirements can be met from solar power. One set of experts said at the most 5%. Another set of experts said at the least 85%. So there is this yawning credibility gap.

So there is a problem there. If you are doing cognitive mapping of experts you are going to get these different clumps of experts, and then how do you deal with them?

In the paper about halfway through there is a diagram from John Adams about the risk compensation hypothesis. It is a balancing act between risks and benefits. He has the idea that different people will have a level of risk with which they are comfortable. They have a risk thermostat set at a particular level. That seems to me to be exactly the same thing as risk appetite. Risk appetite is defined as the level of risk that the firm, the organisation, is comfortable with.

But that diagram is only the first stage of the argument of John Adams. It says there is a risk thermostat and firms will have the thermostat set at different levels. But then the big question is: how does the thermostat get set? How is it some firms’ appetites get set at levels quite different from others?

John brings in what is called the cultural theory or the theory of plural rationality. There are cultural filters on some of his arrows. Some risks are magnified. Other risks are put into the background. Some reward benefits are magnified. That is where the contradictory certainties come in.

I would just suggest that if you have this risk uncertainty and certainty, and you added contradictory certainties and then further on in the paper you take on board the cultural theory hypothesis, which is the second part of the argument of John Adams, the cultural filters on his diagram, then you have the whole lot.

Mr M. G. White, F.I.A.: I would like to mention the Kay review of UK Equity Markets and Long-Term Decision Making. The Secretary of State for Business has asked John Kay to examine investment in equity markets and its impact on the long term performance and governance of UK quoted companies. He is looking at the generation of wealth right through to the effect on the underlying saver. There is currently much interest in the question of short-termism. The first call for evidence that they issued came out in September and the submissions in response to that were collected just over a week ago. It is relevant to today's topic. One of John Kay's advisory team is a member of this Institute as well as being a past chairman of the NAFI.

Part of today's brief is to ask how firms should define and use risk appetite. I would have thought that might include asking what the risk appetite should be. The way in which I have interpreted the paper is that we take what the risk appetite is and we then see what we can do to operate in practice in line with that risk appetite. I do not get the impression that there is much focus on asking what the risk appetite should be.

It also implies thinking about whether the company does actually work as a team to deliver on the risk appetite rather than allowing individuals to pursue their own agendas. If we look at what has really gone wrong in the high profile cases, "own agendas" is a rather large part of it. I should like to have seen a bit more discussion on the problem of motivation of individuals. People who have heard me speak before will know I am keen on this.

For a business whose *raison d'être* is that of accepting risks to meet the needs of the customers, today's research could potentially be very useful. The impression, as other people have said, is there is much very good and useful material in the paper.

The Kay review is going to focus quite heavily on the pressures which companies feel from their shareholders and all the behavioural drivers that follow from that. The biggest risk facing a business is its leadership.

A short CII (Chartered Insurance Institute) discussion document arrived in my inbox the other day. Its title was "The Road to Ruin: Next Exit?" Its subtitle was "Insurance Reflections on Corporate Governance and Risk Management". You can go on to the CII website and find this document for yourself. Suffice to say that the conclusions included the following:

- Traditional risk management, and I think today covers that ground, would appear powerless to control many of the potential risks inherent in board level behaviour. Such risks are barely even recognised within the traditional risk management frameworks.
- The importance of leadership personalities, how well the board actually understands the business, and how well it questions the executives. My own thought here would be if the board do not

understand the business well enough they should not be there. In other words, I would almost start there in my risk management.

- The importance of how remuneration and other incentives operate to influence behaviour.

Mr Ryan: There are two other points that I would quickly make. Mr Smith's point is a very valid one. Filling in complex matrices like that is very difficult. The only approach I have had is to try and do some of the computations myself and then try to translate the results into a format that the individual might be able to understand and try to get them to do "yes/no" answers and then draw conclusions from that. It is not terribly effective. So if anybody has done anything else in that area it would be useful to hear about it.

The other thing we as actuaries sometimes tend to forget when dealing with these operation type risks is central limit theory can apply. There are quite a number of risks and they are similar even if they are all skew distributions.

Mr R. J. Houlston, F.I.A.: I found the paper, although it was titled risk appetite, concentrated on risk management. The issue of setting risk appetite did not really get covered to any great extent although in Section 6 under 'Step 1 – define objectives' there are five principles which seem like a very good starting point when setting risk appetite. In this way risk appetite relates to the outcomes for the firm's balance sheet or operational issues.

I would be interested in knowing whether other people believe we should have a better set of definitions for terms used in risk management. I would like to see some simple standard descriptions for risk appetite, risk tolerances, risk management etc. Either that or could someone please tell me where this has already been done?

Mr P. H. Simpson, F.I.A.: This is a more a question than an observation. I was trying to work out where the boundaries were between what Mr Cattle termed the 'Donald Rumsfeld syndrome' and just pure ignorance in the cognitive mapping.

I can see the approach is very good for understanding the relationships between individuals. But there may be problems if the individuals involved are not experts in their field and therefore are just missing things. They may not be aware of things although they may be knowable, so not a genuine black swan. If things are knowable but unknown there is a competence issue. There was some earlier discussion on board personalities and the competence of non-executive directors, which I think you could trickle down into management as well.

If information is not known in the individuals you are talking to, how would the cognitive mapping capture that? It strikes me that the limitation of knowledge is probably one of the major risks.

Mr G. D. Clay, F.I.A.: One thought comes to me as a development from Mr Ryan's earlier comment about computational speed. It is a challenge to get any model tractable and therefore you tend to oversimplify. Referring specifically to banking, I wonder if something similar has happened with risk appetite which tends to be seen as a rather one dimensional thing.

There is merit in a slightly different approach, which tries to define various levels of consequence to which you are averse. For example, the first level may relate to the dividend, either to have to cut it or to fail to meet the 5% p.a. target increase that you may have set yourself and told the market

about. A possible second level is being unable to raise the additional equity capital you need after a much bigger setback, which may lead to forced asset sales and a change of senior management. The third level is going insolvent.

My totally non-scientific observation of what has been going on in the world during my career is that most companies have very little equity capital now relative to the norm when I took the exams 40 years ago. Then companies were very careful how much debt they took on or, perhaps more accurately, the lenders were very careful how much debt they would let them take on, because they wanted it to continue to be serviced even after a significantly adverse event.

Over, say, the last 20 years the concentration on maximising earnings per share has driven the risk levels up substantially in a holistic sense, even though the perceived risks appear to have come down, because investors think that since they are modelled extensively, the company can measure them accurately and then manage them effectively.

Is there merit in trying to define risk appetite as a multi-dimensional entity, ensuring that at each aversion level any modelling applies different parameter sets that include much greater assumed correlations as the impact on the company rises?

Mr P. J. Nowell, F.I.A.: The whole concept is very interesting. To be able to understand a business and to be able to structure it so that you can see what drives it is an extremely valuable thing. It is something that should lead organisations to be extraordinarily successful for very long periods of time. In some senses some companies have done that by being perhaps very focused on knowing their business well and just pursuing profit and continuity and keeping a good reputation.

I was struck by a comment about the ‘way we do things around here’ in a business. There are characteristics of the people in the business which are as they are and have grown up through time. The civil service is an example of an organization where this phenomena is common. If the ‘way we do things around here’ is not actually successful, or is a bit getting out of date, it is a good thing to change it. However, there are dangers in somebody new coming in or some system coming in which is different but is not successful or does not take the organisation with it. Therefore that seems to me to be one of the major risks. You can see some very successful companies who have been around for a long time and suddenly they decide to change the way that they do things entirely. GEC was probably a classic.

That is a very important part of the whole thing. This led me on to thinking how do you actually make sure that, once you have developed this modelling and understanding, you keep the show on the road? If you are going to have this model which gives you insights, you have actually got to get the rest of the Board, and/or, the rest of the management, onside with you and keep them onside.

The behavioural characteristic of businesses of one sort or another is that once people understand how the system works, the ambitious ones wanting to make their name or increase the size of their empire, use their understanding to benefit themselves, via remuneration, promotion etc.

I should, therefore, be interested in your view about once you embed this methodology in the organisation, how do you make sure that it is bought into by everybody and retains this buy in?

Dr B. E. Malyon, F.I.A.: I would like to pick up on the cognitive mapping technique used. Cognitive mapping was developed by Colin Eden as a tool to help companies explore and

understand complex decisions or problems often where different people in the organisation have different perspectives or views about what the problem issues really are. He and Fran Ackermann developed a number of approaches for facilitating groups and managing the group dynamic to ensure that a wide range of perceptions and views are fully considered.

Mr Smith noted earlier that we are currently in a euro debt crisis which few people predicted. Cognitive mapping can add value to the process of identifying risks and exploring risk appetite precisely because it seeks to ensure a broad spectrum of perspectives, views and scenarios are considered. Furthermore, the cognitive mapping technique is designed to encourage buy-in from participants and a shared view or perception of the problem. The paper focused on the outputs from the case study, but did not discuss the process. I would be particularly interested to know how the process worked in practice and how it could be implemented by insurance companies to help explore and understand risk appetite and embed a risk aware culture within the organisation. I would suggest that this could be a fruitful area for future research.

Mr Houlston: This may be slightly off the topic of the paper, but one of the areas I considered for systemic and evolutionary risks was how do we spot the next big risk? The analysis in this paper is talking about the evolution of things that are already there. I asked myself could we find the next new problem like mis-selling or Northern Rock.

One idea I had was to analyse rates of change. For example, in the personal pensions boom from 1988 there was a huge increase in the sale of personal pension plans. Maybe operational risks crystallised as a result. It is just an idea that we ought to be looking for other risk indicators. Whether rate of change is a good indicator or not I don't know.

There was a comment about the European crisis. I suggest going into a single currency was a massive one-off change. Any time there is a big change, perhaps we ought to set a little warning flag in our mind because we can never know exactly what is going to happen.

Mr P. D. Needleman, F.I.A.: In answer to the question that has just been posed, if you look back at the emergence of major risks, they are usually known about for some time. In fact the recognition of emerging risks often follows a common pattern. At first, they are not widely recognised, then there is a period of 'self-denial' by those likely to be most affected. Only after there is media attention and broader public recognition do emerging risks get accepted and addressed.

If you take the emergence of smoking as a major risk over the 20th-century, there is plenty of evidence to show that the way that people reacted to the issue, particularly those with vested interests, actually made the risks much worse in the end. There was a long period of denial before people began to listen to those who had been talking about these risks. You could say that mis-selling in the insurance industry had a similar development pattern.

How we respond to emerging risks is critical. It depends on the culture in our business and whether we are prepared to listen to people who are providing an alternative view before it becomes the accepted wisdom. If so, then its more likely that we can spot emerging risks early on.

Mr White: On that topic, I think the source of some of the major risks is that wishful thinking keeps them suppressed for too long. An example might be the belief that economic growth will get

us out of trouble, whereas we ought to realise that we are not individually competitive in the world and so much further grief has to happen. It is just at what point will we admit it?

Professor N. Allan (responding as a panel member): There was a question from Dr Pryor about whether we look at risk from an evolutionary perspective as a species? It is a good question. We have included an explanation of this in tables 7.1 & 7.2 in the paper which tries to convey how indeed risk can be thought of as a living system. Work by Pagel in *Nature* argues this same question relating to language. You might argue that it is not a species, but it does indeed appear to evolve in a similar fashion. Others such as McCarty have done the same for manufacturing companies. Certainly many social scientist would argue that organisations evolve in a very organic way and Morgan uses it as a metaphor in his organisational theory. I would argue that risk is a social construct anyway, so it is not a huge step in my mind to viewing risk as a socially evolving and adapting system. Some of the techniques that have been used such as phylogenetics have been successfully applied in biology down to the bacteria and virus level, they are probably also not species as we would understand Darwin's approach to it. Without going into the detail, I think that there are strong analogies between what we can think of as being a living system and risk. A lot of people here are making money out of this thing called risk. It seems to be very real from that perspective. We have some evidence that risks do evolve, mutate and some also seem to become extinct.

Moving on, maybe I can elaborate on another point. Sometimes the data that we have may not be sufficient to say that there is enough evidence in this data for us to project an evolutionary story or a path dependency. Of course not every set of data would produce a valid tree but there are couple of metrics in the paper, that help to assess whether this is the case or not.

I like the idea of a risk called Pat! We would not just call it Pat, we would try to understand what Pat was. As soon as we start trying to identify the characteristics of Pat we might get some sense of it. We might call a bear a polar bear or a brown bear. It is only when we actually start defining it that we find language is important to define these things.

Professor P. Godfrey, Hon F.I.A. (responding): There are a whole load of questions which can be classified under the heading of what is known as epistemic risk. That is risk which is related to our state of knowledge. Rumsfeld is a classic example of that. The corollary of that is that in order to deal with the risk, we have to have the right understanding or right state of knowledge to deal with it.

It is in my view axiomatic that if you take a monocultural view of any complex problem, you are unlikely to have a good understanding of it. What that then means is that you need a diversity of views of that problem. One of the key limitations of most monocultural views is a denial of ignorance. It is almost motivated into experts: they are called experts so you have to provide an answer of some sort. Very incomplete knowledge but an answer is expected. Or, if you take the Japanese point, and the university has done some analysis of this, the actual cultural relationship between the boss and the levels of hierarchy within the structure means that upward criticism is just not on. This can lead to really serious epistemic problems going on. How does that relate to all this? First of all, no one tool is likely to give you a complete understanding of your problem. In my experience, you have to explore the problem with a variety of tools.

One of the advantages of the concept mapping tool is that it is enormously flexible. I have seen it used by people in the Open University to completely map debates in the House of Commons, for

example, illustrating lots of points of view, and bringing understanding of what is going on to a level which probably none of the people there actually realised, which is even more interesting. So I really do think that it is creating understanding that is important. In complex situations emergence comes from the relationships between the parts. So ways of mapping relationships are actually very important to understanding emergence. If we are going to see new emergence, the only way we can do it is by understanding what is going on and being able to reflect that intelligently. So the people involved, their intelligence and their understanding, all become part of the system. No one form of calculation will give you a complete answer.

Mr White: In response to Professor Godfrey's point, I completely agree with what he is saying. I like investing in small companies and getting to know the management. That gives me an edge. What I want more than anything else, the No 1 characteristic, is humility in the board. That means they have to be confident enough to admit they might make mistakes. It is all about the culture, which means a pleasant "family" style of culture within a business if it is going to work properly.

I was talking to a colleague the other day. He recounted an experience when, as an underwriter, he made some significant mistake. He pointed it out the boss with some trepidation. The boss had no problem at all and was very encouraging and supportive. That cemented a strong and enduring trust and loyalty between them.

I think if we focus on issues such as how nice a place it is to work rather than how aggressive, and how much "I am out for me" is encouraged or discouraged, that is going to give important pointers to whether a disaster may be happening.

Mr M. Arnold, F.I.A.: These days I tend to be on the receiving end of this sort of information, chairing the risk committee of a board. There is a lot of work to be done in this area and on how you present this information. Certainly those non-executive director colleagues who sit on the committee with me would struggle enormously with a large part of this. As a profession we need to do a lot more work on how we present the output of this work in a form that lay people are going to be able to use sensibly.

The Chairman: Mr Cantle, your colleagues have responded to one or two questions. I will give you a similar opportunity.

Mr Cantle: There have been a lot of themes around explaining more of the process. The report was long enough and we did not want to give a blow by blow account of everything.

The process of going through the cognitive mapping and the related aspects of the process was a journey in itself for the case study companies. They really appreciated having the opportunity to find out what their colleagues knew, and actually putting the knowledge together. Partly to reply to Mr Simpson's question, "how do you know what you do not know", an incomplete story is so stark on a cognitive map it leaves a great hole in it that it is actually a nice way of visualising that gap in understanding and then exploring that.

Certainly we found going through the journey with the organisations involved in this that it was quite a nice learning experience for them. On the point of multi-dimensional risk appetite, the intention of the model that we presented was to show that it is possible almost to arbitrate between multiple competing objectives in setting risk appetite. A lot of people focus on how big the balance

sheet is and that is the end of it, whereas the framework can actually resolve the tension between non-financial and financial and other various forms of risk appetite statements. The intention was very much to be inclusive of capturing that multi-dimensional aspect.

A lot of the comments seem to be emphasising understanding as being key. I would hope that from the detail in the research paper you can see that understanding is key to what we have tried to capture in the tools and techniques put forward.

The aim is that people such as yourselves can leverage the understanding in the business, capture it and structure it. Then really make sense of it, challenge it and then play that back also in a way that makes sense to people. The language is sensible. It is not purely mathematical or purely analytical. It is something which any participant in the business could understand.

I do not know whether Mr Allan, on the cognitive mapping side, wishes to say anything.

Professor N. Allan: There is a process that we have developed to guide these workshops in how to develop understanding using a cognitive mapping approach. Often the language that is used to initiate the process is quite critical. Mr Arnold's point is well taken and maybe we could expand on that part of the process. The sources that we have referenced are pretty good guidance not only for how to use the tool but for how to set up the process.

In practice this process, is not for everybody. Maybe not every actuary is cut out for running that sort of workshop. I hope that is not being too blunt. This is a different skillset. I am an engineer and Professor Godfrey is an engineer and we seem to have made the transition without too many scars. It is a valid point that this is something that is a very different process to what you are used to.

Before launching in with the board or any executive team, this needs to be piloted and it needs to be practised. The manner you initiate people into a concept mapping exercise, how you capture information and how you actually map it, are not hugely difficult but they are different and the differences are not trivial. It is something that, with a little bit of practice and sitting in on some sessions, is easily achievable. It is a different way of thinking and connecting ideas explicitly. I am happy after the meeting to go through that process in some detail.

Mr T. J. Llanwarne, F.I.A.: I should like to ask a few questions. The first question is: has your research ended or have you got more to do? The second group of questions relates to the success of the research of which I am a big fan. I am interested in questions about outcomes and maybe risks of success or not. What would you see as success in terms of the research you have been doing and the paper which you have written? What would you see as next steps? Are there any preconditions to making a success?

Your research addressed two particular questions. There could have been lots of other types of questions raised. What other questions do you think could have been asked for which your techniques might well have been useful in the field of risk?

Mr Lewin: An example of a complex system is Network Rail or its predecessor, Railtrack. One of the significant events in the history of Railtrack, which led to its demise, was the derailment at Hatfield. It turned out that the rail had been defective for months and the spare rail was lying beside the track in that place, yet no speed limit had been imposed.

If you think about that risk and how it might have been approached beforehand, I believe a concept map would certainly have put “avoiding a major derailment” as one of the main objectives of the organisation. One can imagine that, going back through the process of concept mapping, one would have been able to come up with the management action that you needed to impose speed limits when the rails were defective. That is obvious.

What you would not be able to do by that process, however, is to spot the particular case of that particular rail which was lying beside that particular track and the fact that no speed limit had been imposed. You might have processes which should have picked that up. But if not, you would then be relying on your local staff to report that that risk existed.

I think it illustrates very nicely the need to bring in a good risk-aware culture, with encouragement of all staff to report when risks exist, no matter how much concept mapping you might have done beforehand.

The other thing which contributed to Railtrack’s demise out of that incident was the way that they handled the risk after the event. The way they handled the crisis. You will remember that they imposed many speed limits all over the country and caused a lot of unnecessary inconvenience to people. The way that was done was viewed very negatively. So you had a cascade of things going wrong. First of all, the rail not being replaced and no speed limit; then, secondly, the accident; thirdly, the way the risk resulting crisis was handled.

Some of the techniques outlined in the paper, of which concept mapping is probably the most important, point to ways in which proper procedures can be devised to try to minimise such risks for the future in any large and complex organisation.

Mr Ryan: I am just going to come back on the points made by Mr Lewin about not handling the risk properly, which, in a sense, is outside the scope of this paper. A major aspect of risk management is control of risk, which this paper is not, quite rightly, covering at all. Disaster planning, and so on, is a major part of that.

I think that there is a lot of useful stuff that can come out of this that can go into those types of plans. It is very important that we separate out what we are trying to do here, which is quantify risk and identify risk. That is only one quadrant of the circle. But it is very important and we should not underestimate that. It is also very important that we recognise that this is not the whole solution. This is quantification and identification, which are very important but not everything.

Mr E. M. Varnell, F.I.A.: The thing I particularly like about the paper is it is really taking us away from the reductionist approach to economic capital and ERM, which has certainly plagued the banking industry and led to some of the problems we saw in 2008 and has been the de facto method that has tended to be used in the insurance sector.

It is just now, when people are trying to make sense of a lot of those massive models that they have built for themselves that we are actually seeing, if you like, a crisis of the reductionist approach. It is a major challenge to get non-technical people, non-executive directors, for example, and people on boards who are not familiar with what a T-copula or a Gumbel copula or a Pearson IV distribution, to understand what risks their businesses are really running, and how they can mitigate those.

This even occurs when we start to think about some of the risks which we think we can quantify. So, for example, I have been involved with Mr Cantle on a working party looking at extreme events. The best that we can really do there is to apply statistical techniques to historical data.

The key assumption that we always make when we look at historical data is that the system that created that data is the same throughout the observation period that we used. Whether or not we use data going back to 1900, for example, to try to estimate the one in 200 equity fall, it is a very strong assumption that the world was wired up in exactly the same way throughout that period.

My favourite analogy is the Danish stock market. Denmark had a relatively nice little farm-based economy for quite a while. Relatively recently they started to get a lot of mortgage-backed securities in that market. Before 2008 the largest ever loss on the Danish stock market had been 22%. In the crisis year of 2008 the market loss was approaching 50%. Nothing in the historical data set could possibly have told you that that was going to hit Denmark.

Of course regulators are now starting to pick up on this, realising that it was actually a contagion and a systemic problem that caused 2008. They have started to use a systems approach in their thinking. There are some references, if you want to look them up. Andy Haldane has been very prominent in trying to bring systems thinking into the Bank of England in their regulation of the banking sector in the UK. Some work has also been done on using a systems approach at the European Central Bank.

There has also been work in industry. A couple of years ago in Edinburgh at the ERM conference we had Riccardo Rebonato, who has been a big supporter of bringing systems thinking and Bayesian networks into enterprise risk management through his work at RBS.

Mr Llanwarne has challenged us to think about ways in which we could use the next example. Mr Thompson of the International Institute for Applied Systems Analysis in Austria, who is one of our VIP guests, has been doing some work with Dave Ingram of Willis Re and of the International Actuarial Association. They have been looking at how you can use some of those cultural theory constructs, which were discussed earlier, in order to understand how we can have relatively chaotic development in economic systems. If people have not read those articles that Mr Thompson and Mr Ingram have been producing, then I can heartily recommend them. They make for a very interesting read.

Finally, a question back to the panel. When we were up in the Redington Room, we were having a very interesting discussion regarding financial services just starting to pick up on systems thinking and that this is also emerging in other areas such as civil and more general engineering. So I would be very interested to hear from Mr Cantle and Professor Godfrey about how other industries are starting to adopt a systems thinking approach as well.

Professor Allan: It is certainly not the end of the road. I am engrossed in this topic. We are only scratching on the surface of applications from complexity science and systems science that can be applied to many disciplines, including yours.

What was particularly successful, and quite surprisingly successful, was the openness of the pilot companies that allowed us to use their data and the openness with which they accepted and took on

board the results. To be honest, I really was not expecting that. That seemed to be a great success. I think it bodes well for some of these approaches and their adoption.

Going forwards, some of the things I would have liked to have done is apply the techniques to other areas. For example something we have just touched on, are losses. This gets away from “is it a real risk or not?” If we have real losses, the evolutionary approach is particularly useful because it becomes more tangible.

I look forward to getting invited by many of you to investigate your data.

Professor Godfrey: The first question was: has the research ended? I am going to give a systems thinking answer to that. Basically, for any complex issue, and risk issues fall into that area, they cannot be solved. Because of the nature of the complexity, they can only be resolved. So, a point I shall be making later is that we have to be humble. It was a point made in the audience as well.

We have to be humble about what we are doing. We can all learn, and we can learn a lot better if we learn together. That collaborative learning culture actually underlies a great deal of what is necessary in this area.

The second question was: how would you see a successful outcome? First of all, I would say that I think the openness of the questions and the discussion that has been going on here is an indication of a successful outcome emerging.

A successful outcome would be making a difference by recognising the incubating factors that actually lead to the risks that we are talking about, and some of the big ones that we have talked about here, like the financial crisis.

There are plenty of people who were saying that bubble has burst. Experience has shown time and time again that when things get to levels of absurdity in terms of the rational drivers of things, something is going to happen. Yet, as has been pointed out, there is a great deal of motivation not to sink the ship – not to rock the boat may be a better phrase.

So if this allows us to make a difference, allows us to listen to those incubating things that are going on, and at least do something about it, use a little bit more wisdom and caution maybe, then I would feel that we are getting somewhere and the tools help in that respect.

What other questions could have been asked? I would say you could have asked the question: how can companies use these techniques to recognise their epistemic risks? I think that is a very important and valuable question, and there are very valid answers to be drawn from that.

Mr Cantle: I am in agreement with the other authors. I have no need to repeat everything that they have just said. On behalf of the two authors who are here and the one who is not, thank you very much to the meeting for such a thoughtful discussion. It is an interesting and new topic and it is nice to see everybody engaged in talking about it in an open way.

Regarding Mr Llanwarne’s question, success for us is to bring this sort of thinking to the meeting and have it augment the toolkit which we currently have. It is meant to be something which adds on

to the body of knowledge which we have developed as actuaries in the past and give us something additional to put in the toolkit, as it were.

We have used these techniques over many years in a very wide range of ERM and business applications. We hope that today's meeting will give you the confidence to do some of that as well. Thank you very much.

The Chairman: May I thank Mr Cantle and his fellow authors and can we all join together in the traditional way to say thank you?