

ABSTRACT OF THE EDINBURGH DISCUSSION

Operational risk dependencies

[Institute and Faculty of Actuaries, Sessional Event, Edinburgh, 20 January 2020]

The Chairman (Mr D. J. Grenham, F.I.A.): My name is Dermot Grenham and I am the leader of the Scottish Board, and it is my pleasure to introduce our speaker Patrick Kelliher.

Patrick is a Fellow of the Institute and Faculty of Actuaries (IFoA) based in Edinburgh. He is managing director of Crystal Risk Consulting, an independent risk and actuarial consultancy. He has extensive industry experience of operational risk modelling. He is chair of the Operational Risk Working Party and is a member of a number of other actuarial professional risk management working parties.

Mr P. O. J. Kelliher, F.I.A.: Before the discussion, I thought I would give a brief synopsis of the paper.

The first thing that I want to cover is: why would we want to model operational risk dependencies? From a risk management perspective, I would always want to know how my operational risks interact with each other to understand the potential for losses to accumulate over different categories, but also to understand the diversification between operational risks and between operational and non-operational risks, such as market, credit and insurance.

I think that for any model of risk based capital requirements, it is important to allow for such diversification benefits to give a true picture of exposure. In terms of the regulatory background, with Solvency II, you have two kinds of approaches: internal model and standard formula. For internal model firms, they can set their Pillar I regulatory capital requirements in relation to internal models of risk. I would expect the internal models of operational risk they use to allow for diversification, not just between operational risks but also with non-operational risks.

With standard formula firms, the standard formula does not allow for any diversification. As part of the Pillar II assessment of the appropriateness of that standard formula charge, you may wish to consider diversification between operational risks and non-operational risks.

In terms of banks, I think they are on a different journey from us. Under Basel II, you had the advanced measurement approach whereby the banks could use their own models of operational risk for Pillar I regulatory capital. As part of that, they could allow for diversification between operational risk, but, crucially, not between operational and market or credit risk.

Unfortunately, banking is moving away from internal models, and the advanced measurement approach, where we had internal models for operational risk, has been replaced. Now everybody is on a standardised measurement approach, which would be quite similar to the standard formula for life insurers.

However, banks will still probably want to model operational risk as part of their Pillar II internal capital adequacy assessment process (ICAAP). As part of that, they would probably allow for diversification between operational risks, though my understanding is that it is very rare for them to allow for any diversification benefits with credit and market risks.

Asset managers are in a very similar position to banks in terms of wanting to assess operational risk as part of the ICAAP. For asset managers, they tend to have very modest non-operational risk

exposures, and operational risks are probably the key exposure they face. They will probably want to allow for diversification between operational risks in that ICAAP.

Turning to the modelling of dependencies between operational risks, and diversification between operational risks, the first thing we note is how diverse is the category of operational risk. It covers everything from conduct risk to cyber, to processing errors, to business continuity. Under Basel, you have 7 high level I and 20 level II categories of risk. I was part of a working party to identify another 300 subtypes of operational risk. There are a lot of different types of risks under the one heading of operational risk.

It is obviously implausible that we are going to see all of those crystallising together. I think that it would be appropriate to allow for diversification between operational risks, but how much? The challenge is that a lot of times operational risk categories may seem uncorrelated at first glance, but, if you dig deeper, there are often some common drivers which point to a dependency between operational losses.

The working party identified a number of drivers for losses, including people – relating to weaknesses in recruitment, training and retention. For instance, recruitment could result in hiring the wrong kind of people, which could result in fraud or processing errors.

Model governance is another driver. Weak model governance could lead to flaws with product pricing and could also lead to financial reporting flaws in terms of model valuation of assets and liabilities.

The compliance culture is obviously important. A weak compliance culture could lead to losses arising across a wide range of conduct risk categories. Then there is the overall driver of weak governance. If you have weak governance overall, then you could see losses cropping up in very diverse categories including fraud, conduct and business continuity, if you do not have the business continuity planning in place to be resilient to events.

To give you an example from the paper, which is that change management is an important driver of operational losses. If we take the example of change management and poor project management practices, you could have the situation of flawed implementation with poor controls which could expose your customers to fraud and cyber risks. You could have document production errors which could give rise to conduct failings, but also could give rise to product risks and to inadvertent guarantees in your documents.

There could be processing errors or financial reporting problems. In my experience, financial reporting is generally the last thing to be built in terms of any systems. Obviously, any weakness there in terms of what is built could emerge in future years as financial reporting and management information (MI) problems. You could end up with a very unstable system which is much more prone to outages.

You can see that just one linkage point of weak project management or change management skills could affect a wide range of operational categories. Just to reiterate, even if they do seem, on first glance, unconnected, when you look at these main drivers there can be a lot of connection.

In terms of coming up with the correlation assumptions, the starting point is going to be empirical loss analysis, but the key problem for most insurers is lack of data, particularly of the low frequency/high-impact events. This means that your data could be missing a correlation because it is not long enough to cover the emergence of two extreme events together, or alternatively you could sometimes have spurious correlations.

I have come across an example of where a conduct risk happened to coincide with an IT outage. There was no actual link between the two, but because it happened in the same quarter that led to a very high correlation estimate. It is a bit of a problem with lost data that you can obtain some very peculiar results.

Another issue with operational risk is that the frequency of events is often quite low. That can result in systematic understatement of the empirical correlations.

Just to elaborate on that, the working party did a very crude simulation exercise where we had five operational risks with models with random variables between zero and one. If the random

variable was less than 0.1 a loss was assumed to occur, otherwise no loss was assumed. We aggregated those assuming a Gaussian copula and 50% correlation between these variables. We found that the actual empirical correlation between the loss events was only 25%.

That is an example where you can obtain systematic understatement of correlations when you have very low frequencies.

One area that can help in terms of loss data is industry studies. From the studies that I have seen, they generally point to correlations of 25% or less. One paper that we did identify as part of the literature review was a 2007 paper on US banking losses. That pointed to correlations of 30–50% which was considerably higher than in other papers.

I think, ultimately, empirical loss data is a starting point. We also need to supplement that data with an expert judgement.

The problem that we have with expert judgement is the sheer number of assumptions that we need to handle. If we have 20 operation categories trying to populate a 20 by 20 correlation matrix, we need 190 correlation assumptions. It is going to be very challenging, to say the least, to ensure that those are all set with appropriate rigour, reviewed and challenged, and meet the positive semi-definite criteria necessary for correlation matrices.

We have identified two possible solutions to this problem. The first one is to group risks. Rather than having a 20 by 20 matrix, you would split them into four groups of five. You only need 46 correlation assumptions which might be more manageable.

In terms of how we group risks, there are a number of different suggestions. One would be to use level one categories in Basel II. Another one would be to group by functions. You could map certain operational risk types to functions. For instance, mis-selling and other conduct risks could be mapped to sales and marketing, employee relationship risks to HR, reporting errors to finance, and so on.

Then there are the general generic types of people processes, systems and external events that form part of the Basel and Solvency II definition.

Another potential way of addressing and setting correlations, and bringing expert judgement to bear, would be to consider the impact of a suite of scenarios on losses in different categories. For instance, we might assume a flu pandemic creating business continuity losses. When we consider that, we might also think that that is likely to affect processing errors given higher levels of death claims. That would point to a *prima facie* link between processing risk and business continuity risk.

Similarly, a stock market fall could cause mis-selling of investment products, but could also bring to light product flaws, such as inadvertent guarantees given in literature, which could crystallise in depressed markets.

Once we have our correlation assumptions, how we model them? The current practice, among UK life insurers at least, is to use copula aggregation to model dependencies. Gaussian copula is the most common choice.

The reason for the popularity of the Gaussian copula is that they produce a full distribution of operational losses which is required by Solvency II. Under Solvency II, there is a delegated regulation which requires you to produce a full distribution of own funds.

Another reason why we might want to use copulas is, if you look at the alternative correlation matrix aggregation, typically that assumes that the underlying loss distributions are elliptically distributed, which is rarely the case for operational risk and can result in a distorted picture of diversification. That is another reason why UK insurers, at least, are using copula aggregation of operational risks.

In terms of the choice, obviously Gaussian copula is easier to implement and requires fewer assumptions. We could use a T-copula, but we need to make a further, probably subjective, assumption with regard to the degree of freedom parameter.

One problem with the Gaussian copula is zero coefficient of tail dependence, but I do not see that as a huge problem for operational risk aggregation.

If we look at the 99.5th and the other percentiles that drive our economic capital, the Gaussian copula can model extreme events happening together, depending on the correlation assumptions. Some very crude simulations of conditional probability: for example, the probability of a one-in-200 loss under Risk B given a similarly extreme loss under Risk A, assuming a 25% correlation, a Gaussian copula gives you a conditional probability of 2.5%.

That might seem small, but it is five times more than the independent probability. Assume a 75% correlation, and the probability of an equally extreme one-in-200 event, it rises to 27%. You have a one-in-four chance of a one-in-200 event under B given a similar event under A.

This highlights the sensitivity of dependency results to the correlation assumptions. Those correlation assumptions are likely to be very subjective. The feeling of the working party generally was the more sophisticated copulas than Gaussian are probably somewhat spurious given the subjectivity of the correlation assumptions.

I should like to turn now to cover dependencies with non-operational risks and the diversification between operational risks and market, credit and insurance risk.

The first thing I would note is the degree of asymmetry. Market, credit and insurance events drive operational loss, but the reverse is rarely true. The one exception I would highlight is reputation damage: I see that as a vector of transmission between operational losses and insurance. You obviously need a collapse in sales leading to higher lapses. Both of those would drive up unit costs, so we can see there how reputation damage transfers the operational loss and impacts on the lapse and the expense elements of insurance risk.

Another area to be aware of in setting dependencies between operational and non-operational risks, particularly for insurance risk, is implicit allowance in non-operational risks for operational risk events.

Just to elaborate, take insurance risk. A lot of times the models are based on historic claims analysis. Historic claims experience will reflect not just claims, but also might reflect underwriting and claims processing errors and fraud to the extent that they are not identified and stripped out of the experience. There will be an element of allowance in insurance risk models for underwriting and claims operational risks. It is unfortunately very difficult, by its very nature, to say how much, but it is something to bear in mind when considering correlations.

Another aspect to consider is the degree of conditionality; the severity of operational loss is often a function of market and other non-operational risk events. The classic case of that is mortgage endowment mis-selling, where what you saw was whenever there was a fall in markets and fund values that led to higher compensation payments. We can look to mis-selling as essentially a put option written by the life insurance provider, the cost of which will depend on markets.

The final thing I should like to say in terms of dependencies with non-operational risks is the time-lag element. There can be a significant lag between a loss being incurred and the actual loss arising. That is quite important. When we look at it over a long time, we can see linkages between operational and other risks.

Take, for example, Payment Protection Insurance (PPI) mis-selling. We argue that there were linkages between PPI mis-selling and aggressive mortgage lending practices in the run up to the financial crisis 2007–2009. But if we consider the market losses that arose, they all crystallised around 2007–2008/early 2009 with the mortgage losses probably a bit later. But it is only from 2011 that you start to see significant provision being set aside for PPI mis-selling. Most of the cases had not been settled until very recently.

I think it is important to bear in mind you can certainly make the case in the long term for correlation linkages between operational and non-operational risks. When you are looking at a 1-year time frame, you are going to be looking at market and credit losses and insurance losses arising in the coming year, with operational losses crystallising, but which may have been incurred 5–10 years ago.

That is something to bear in mind when you are setting your correlation assumptions. There is a kind of disconnect between the losses you are going to face in the coming year and the current market risk exposures.

The Chairman: Thinking about the different organisations that you were talking about: banks, insurance companies and asset managers, I presume that the relative operational risk importance varies?

For asset managers it is quite important. When you come to insurance companies, is it not that material? Or is it always something on which you need to focus?

Mr Kelliher: I think operational risk is one of those risks which has been latent. It has always been quite significant for life insurers, but I do not think that we fully appreciated the true extent of it until we started suffering multi-billion pound losses as a result of pensions mis-selling and endowment mis-selling.

The life insurance industry has taken steps to reduce mis-selling risks. Other risks keep coming up. Operational risk is like a balloon. You squeeze one end and then another end comes out.

For instance, you might remove mis-selling risk by appointed representatives of banks, but then you might have a wrap platform and you expose yourself to a different suite of risks, including client funds and so forth.

Certainly for life insurance companies, it is significant. For asset managers, it is probably their main risk. The other risks are not as material.

For banks, again if you look at the cost, I think the cost of PPI is tens of billions of pounds, which I think is on a par with the credit losses that they had during the financial crisis.

It is important for all financial institutions. It is just the degree to which other risks might offset that and, if you are holding capital for operational risks, credit risk and market risk, to what extent diversification is allowed between the three.

Mr J. E. Gill, F.F.A.: One of the things that stood out for me in this paper is a lot of work on justifying correlations, and on making sure we understand them. What evidence did you see that all this work on understanding correlations and so on, is leading to better management of the underlying risks?

I do have a concern that while a lot of work goes into an academic exercise, I am not sure whether the organisations are learning and becoming better at managing the underlying risk itself.

Mr Kelliher: I agree with that. There is a lot of good analysis about operational risk in general. Whenever you are trying to assess what your operational risks are, discussions can be very fruitful. But as you said, particularly for the correlations, to what extent is that feeding into trying to improve poor governance or poor change management practices.

We might obtain those correlations. We might say we have an issue here, but to what extent is that getting through to management? I am afraid it probably is not, which is a pity. For me, the operational risk assessment, the individual marginal distributions and the aggregation, is a useful exercise in itself regardless of the capital figure that comes out.

For me, the real value is the discussions around the scenarios and how the scenarios might be linked, and how that can feed through. I have had some interesting discussions on scenarios that feed into operational risk. The discussions on correlation assumptions I do not think have been fed into business as usual risk management, as I think it should.

Mr A. R. Wallis: I am not sure whether the regulators have a role to play, but is there a danger in the effort put into identifying the correlations that you increase your ICAAP capital holdings required? If so, is there then the potential that the institutions do not invest in the efforts to identify these correlations?

I am thinking in terms of policing that consequence and stopping such practice.

Mr Kelliher: I think that there is a general issue with operational risk that there is a very large degree of subjectivity with assessments. We cannot escape the fact that management might have a certain preconceived idea of how much operational risk should add to the overall economic capital requirement.

If it does push higher in terms of the operational risk capital requirement, then we could push back. A concern that I have about operational risk is, particularly in the banking environment, of not allowing for diversification of non-operational risks – you either end up with a very high addition to economic capital as a result of operational risk because you do not allow for diversification or, when you come up with an initial value, senior management takes one look and says that the number is too high, and then there is a push to think again.

This is the problem with operational risks. Due to subjectivity, you could be pushed to reconsider scenarios, dumb them down and end up with the worst of all worlds. In that situation you are lying not just to the regulator but you are lying to yourselves about your exposures.

Take an example of a bank in the run-up to the financial crisis. It may have had, under an internal model, an operational risk capital requirement of around £2.5bn. But the same bank then incurred a £3bn charge on PPI mis-selling.

So the question is: did they properly consider their PPI mis-selling exposure? Could it have been that they allowed for, let us say, £3bn plus all the other items, say £5bn, and then there was push back? We will never know.

Another concern that I have is false prudence. We are not allowing for diversification between operational risk and market credit risk in banking asset management. What happens is the operational risk is pushed back and compromised to fit in with a value that some management members may have in mind.

Mr A. J. Rankine, F.F.A.: First, a general observation that capital is probably not a great mitigant for operational risk as a risk type, and therefore does the additional complexity inherent in most of these approaches really benefit the companies that are implementing it?

Second, picking up on the theme of a couple of the other questions, I am wondering whether, in the light of the banking industry's move towards more of a formula-based approach, management and, potentially, regulators would be better served with a simpler approach which put more emphasis on individual scenarios combined in a simplistic way, but well understood rather than the additional complexity of a copula-based approach.

I am thinking particularly of a slightly crude example whereby you could easily have management assessment of operational risk reducing from period to period because of a change in the assumed correlation between people and process risk rather than, for example, an underlying reduction in either of those risk types.

Mr Kelliher: I think that is a very good point about capital not being a great mitigant of operational risk. The idea is to have proper controls. There will always be some level of residual risk. We need capital to cover that residual risk.

The other problem we have with operational risk is the lag effect. You could do something about controls now. What is going to hit you in the next couple of years is the result of past sins. We do need to have capital.

I agree about the complexity of modelling. I am generally in favour of cheap and cheerful methods of aggregating operational risk. For standard formula firms, or for banks that are doing the Pillar II assessment of operational risk, there might be a case for a simple correlation matrix.

I have been involved with smaller firms using standard formula approaches. I thought the simple correlation matrix had its limitations, but it was cheap and cheerful and very easy to implement. That is why we went with that approach.

That banking is moving to a standard management approach is somewhat lamentable given that they have ditched the advanced measurement approach used in internal models, although there were problems with it and there was bias.

The concern I have is that there is a mechanistic formula for setting Pillar I capital requirements. It does not strike me that one size is going to fit all. It would have been better to allow internal models, but to strengthen the governance around scenarios to offset some of the biases.

On Pillar II, my understanding is that the PRA do look at scenarios as part of their review of ICAAPs. They look at the scenario analysis produced by the firm. My concern is that it is important to have that review and it is also important to have some allowance for diversification between those scenarios rather than simply adding things up. Otherwise, you are not allowing for diversification between operational risks, which is completely excessive and you can end up with push back.

We come back to the problem of bias. If the end result is a figure that is too high for management, the whole scenario analysis process can become compromised rather than being an honest assessment of exposure. You can end up trying to play a game to arrive at some figure to offset the undue prudence that the regulator specified.

The Chairman: Perhaps I could ask the question of the audience. How many people here work or look at operational risk in their organisations? What is your experience? On the ground, what are the things that you find challenging or difficult? I note Alan (Rankine)'s point that it is all too complex.

Mr A. J. Clarkson, F.F.A.: Probably the biggest challenge is persuading the executives and the board to invest the time to understand what assumptions are being made and what is the impact of those assumptions. Going back to the points that John Gill made at the start, it is hard to use the approach in a way that makes any meaningful difference in managing operational risk to help them in running the company. It is very difficult to move it beyond being a theoretical exercise that gives them a number in terms of capital that they have to hold, which is unfortunate.

I would sum up operational risk dependency as something which is not easy and is highly subjective.

That then takes me to something Alan (Rankine) said about having a simplified approach that does not introduce spurious accuracy. It is very difficult to get the executives and the board to understand and form a judgement on appropriate assumptions. It is equally difficult, I would expect, for a regulator. If you were a regulator, how do you ensure consistency between companies, given some of the challenges that were talked about earlier?

Mr Kelliher: I think a regulator has issues with operational risk in general. The banking side has just given up on internal models. They have said: here is a standardised figure based on lots of very complex studies that they have done. They have more or less given up on the whole modelling aspect. Like I said, that is a retrograde step.

I think that the modelling process can be useful, as you said. It depends on the quality of conversations and how far up they go. I have been involved not so much in correlation assumptions but in terms of scenario discussions about cyber exposure in particular. That received a lot of attention. The process highlighted exposures that senior management did not expect. They were good conversations.

Coming back to the regulator point of view, I can see the difficulty in saying internal model firms all have their own models of operational risk. They will all have different assumptions. How do we get some consistency?

The key thing is to look at the process, not so much the results. What is the process for arriving at those results? How robust is that process? What is the quality of the conversations?

There is no fancy maths. It all comes down to minuting of discussions, having the right people in the room, having good quality discussions and making sure that the discussion is fed up to the higher powers that be. That includes the key takeaways from the assessment.

I would want to see evidence of such an approach, if I were the regulator, rather than whether we are using a T-copula.

In terms of complexity, people are saying that the Gaussian copula is very complex. For internal model firms, it addresses a certain requirement that is in the Solvency II regulations. They need to produce a full distribution of operational losses.

The working party was saying you could use T-copulas; you could use a lot more Vine Copulas and other weird and wonderful dependency methods. We do not really see a huge amount of benefit, given the subjectivity.

There is one exception. One of the most complex methods is Bayesian network modelling. I do not think that many people have adopted this approach so far, but it is a kind of Holy Grail coming up with a holistic distribution of losses across all risks. I have always found the prospect very complex.

I am aware of one business unit which bought Milliman's Bayesian network model and parameterised it. They were pleased with the insights it gave in terms of what drives operational loss. It was built on drivers such as people, systems and staff turnover.

I generally agree that complexity is not justified for operational risk models, given the subjectivity. Having said that, this Bayesian network could be quite an interesting approach going forward.

Mr H. R. D. Taylor, F.F.A.: I was thinking back to my practical experience a long time ago running defined contribution (DC) pensions operations and also running back office banking operations for a large UK bank.

A couple of thoughts at a fairly high level: one is around industry structure and the trend, and what that means for operational risk, and the other around people.

The first one: I wondered what your thoughts were on the general trend in the industry of concentration of risk and business into a smaller number of bigger entities, and the mechanism for reducing the operational risk in them being tighter regulation which might, or might not, work.

The two examples that I would give are the growth over the past 10 years of outsourcing of insurance back office operations. All the complicated systems and activities that can go wrong are now concentrated in a relatively small number of very big specialist outsourcers. That seems to me to be a change in the nature of how, as an industry, insurance operational risk is managed.

The other one is the recent trend following the recognition that there were far too many master trusts. We have just gone through a wave where a huge number of master trusts have reduced to about 35, which is a smaller number of bigger entities. Again, in the world of DC pensions, we are seeing a fundamental restructure. There might be some interesting aspects of operational risk to consider.

The second one is around people. To keep it very simple, it seems to me that the amount of reserves that you have to have can be a function of how effectively and quickly you can respond to an operational risk event occurring. In terms of people, that might be around whether the people who are working in operations or the business help prevent risk events happening.

One example I would give, of a change in the last 4 or 5 years, is if you want to make a large cash withdrawal from your bank account, you are likely to be asked by your bank or your building society teller what it is for and whether anyone approached you?

There has been a lot of work done by the banks to try to avoid customers being involved in scams, and I would say that although a scam is something that affects the reputation of the provider and the bank, prevention is a big thing. People are the core of prevention, supported by some systems.

Second is early warning as a concept. Your people can give you early warning. I imagine some of the senior executives at Boeing, particularly within the cadre of senior executives who had to leave the company, wish that they had listened to what some people internally were saying about the way that they were developing and rushing through a particular aeroplane which caused a huge number of deaths.

There is a lot more to come out of that, but it shows the importance of at least having a process for listening to early warnings that you are receiving from your people. The TSB melt down of systems was another example, although there has not been anything made public.

Finally, probably the most important thing is speed of response when a risk event occurs. That can be heavily dependent on what preparations you have put in place beforehand, and also structurally in terms of your software and the way that people are operating the process. It is how easily they can respond to something going down somewhere.

One example I can give you was from a live outsourcer. They had multiple processing sites, one of which was in India. Overnight, one local very senior politician died, which meant the next day that nobody turned up for work because it was a day of mourning. With about 4 or 5 hours' notice, they had a major site with no one at it.

How easy was it then for them to balance the load across their other sites in the UK? The answer was, because of the particular type of software that they were using to manage multiple sites, they were able, seamlessly, to absorb the extra workload. There was no derogation to customer service; there was no increase in complaint levels; there were no FSA-reportable events. That is a perfect example of something which could have been a disaster turning out to be something that was well managed.

So a question about the structure of the industry: is it a good or bad thing to move to a smaller number of bigger entities or does that just concentrate risk? And another about the power of using people with the right processes to be able either to give you prevention of risk events or early warning of risk events when they are in train. But when they hit, how fast and how effectively can you respond to them?

Mr Kelliher: There are a few points here. To take the issue of industry structure, if you look at Basel and the new standardised measurement approach, the premise of that is that larger, more complex banks are more exposed to operational risk, and hence the higher capital charge under the standardised measurement approach.

I am not 100% certain. A small bank might not have had huge operational losses. Is that just good luck or because they are better managed? I feel that the regulators think that the bigger and more complex that banks become, the more likely there is to be operational risk.

You mentioned that there is a particular issue in terms of outsourcing. There are some key outsourcers. I think the FCA mentioned in part of their review last year that they are concerned. I imagine that if Microsoft went down we would all be in trouble. Everybody is increasingly using one or two key providers. In the insurance industry, what would happen if Moody's Analytics suddenly went bust in the morning? What would we do for economic scenario generators?

The thing about operational risk, and what I like about operational risk, is that it is constantly changing.

The concentration and consolidation of banking and insurers was one aspect of operational risk exposure. Another would be the rise of cyber risk. Five or ten years ago it was not huge; now it is massive. Within cyber risk we have seen a move away from data breaches to ransomware becoming increasingly important.

But, again, some risks have gone. Mis-selling risk – once bitten, twice shy. A lot of banks have effectively removed that risk by not offering advice. That is linked to the point about the general democratisation of risk and how we are no longer taking on risk but passing it onto the individual.

It is quite a complex picture. Some risks go, some risks come. New risks seem to come to light. You are trying to mitigate one and some people try to make things more efficient by outsourcing.

But it is like squeezing the balloon. Squeeze one part and then the other end just expands. We then increase our outsourcing exposure.

The point about people is very well made. If you have really good people, backed up and empowered by good systems, they can do a power of good in mitigating risks. Referring back to key drivers, bad people can do a lot of damage. Similarly a bad culture can link between disparate operational risks.

For instance, if your recruitment process is pretty poor, you may have people who do not have the necessary ability to operate complex systems or processes, and you will have a higher incidence of manual processing errors. You may also have higher levels of fraud.

That links in with culture. You can have good people. The culture in Boeing seems to have been that people were seeing the problems, but they did not feel that they could raise them. If we look at any kind of major event that happens, we talk about the “known knowns” and the “unknown unknowns”.

For most of us, what is an “unknown unknown”? Some people will be aware of the problem. For most of us it is that particular risk we have never thought about. Some people out there will be aware, most of the time, of something new.

The classic one was in the film “The Big Short”. There were people who could see the US sub-prime market was disintegrating. But the question was: how do you leverage that?

There are people in any organisation who know if something is suspicious. Getting that knowledge up the line is a huge challenge.

The Chairman: You mentioned that operational risk is always changing. How does that enable you to model it if the future is not going to be similar to the past?

Mr Kelliher: In terms of modelling, you have to look at scenario analysis, and trying to think where are we now? Certainly past data can tell you much, but sometimes it can give you a distorted picture. If you are a life insurer and you look at past data, you would say that there is a huge mis-selling risk. But probably there is not.

You always have to look at scenario analysis. As has been said, you might have removed one risk but, taking the insurance industry, they have added on a lot of complexity in the product in terms of SIPP, wrap and drawdown. That has brought new risks.

The only way is to have a forward-looking scenario analysis approach to see what new can happen and also then to try to understand how this risk could interact with all the others.

The Chairman: I should like you to join me in thanking Patrick (Kelliher) for his presentation and for the full answers to the questions that he gave.