Good practice guide to setting inputs for operational risk models by the operational risk working party

Abstract of the London Discussion

[Institute and Faculty of Actuaries, Sessional Research Event, London, 21 March 2016]

This abstract relates to the following paper: Kelliher P. O. J, Acharyya, M., Couper, A., Grant, K., Maguire, E., Nicholas, P., Smerald, C., Stevenson, D., Thirlwell, J. & Cantle, N. *British Actuarial Journal*. doi: 10.1017/S1357321716000210

The Chairman (Mr M. J. Bowser, F.I.A.): My name is Marcus Bowser and I am the sales and practice leader for Willis Towers Watson UK Life consulting business. I am also chair of the actuarial profession's risk board. I am particularly delighted to be chairing this session as many of the early projects that I managed in the risk management space were around operational risk.

I will now introduce our speakers. Patrick Kelliher is the founder and managing director of Crystal Risk Consulting Ltd., his own consulting firm. As well as chairing the operational risk working party, he has extensive practical experience of operational risk identification, management and modelling from a number of assignments over the years. He is a member of the enterprise risk management research and thought leadership committee under the risk board structure.

Previously, he has worked for most of the companies based in Scotland, including Aegon, Scottish Widows and Standard Life. Certainly through the risk board I know that Patrick is very passionate about operational risk.

Secondly, we have Chris Smerald. Chris is a claims actuary for American International Group's (AIG's) Global Property and Casualty Finance Actuarial Department, helping to apply analytics to improve aggregate reserving and performing special studies on claims issues. He is former chief actuary for Lexington London, which included substantial responsibility for developing their advanced operational risk management framework. He has also held a number of volunteer research and leadership positions within the Institute and Faculty of Actuaries, Casualty Actuarial Society, and Operational Risk Insurance Consortium (ORIC).

Finally we have Andrew Couper. He is the group chief actuary and head of risk for Aspen Insurance Holdings Ltd. In his role as head of risk, Andrew manages the group risk management team that oversees Aspen's approach to identifying, monitoring and managing operational risk. In previous roles for Aspen, Andrew has been a chief risk officer for Aspen Bermuda Ltd., and also the chief risk officer for Aspen's two US regulated entities. He is an actuary by background and has been working in risk management since 2010.

Mr P. O. J. Kelliher, F.I.A. (introducing the paper): Before delving into the topic of inputs to operational risk models, I would like to set the scene by making a few general remarks regarding operational risk and frameworks.

The first comment I would make about operational risk is that it is a very diverse risk category. It covers everything from cybercrime to rogue trading; from mis-selling to employee discrimination claims. It is one of the things that makes operational risk particularly challenging to model.

A second remark I would make is that there is a lot of similarity between operational risk and general insurance risk. There is potential for catastrophic losses. In general insurance you have hurricanes. But you have catastrophic losses from operational risks such as the £16 billion loss that Lloyds Banking Group (LBG) has incurred to date on PPI mis-selling.

Like general insurance, there is an issue with regard to legacy exposures; what general insurers would call incurred but not reported exposure. An example of that are the losses in the last decade from mortgage endowment mis-selling in respect of endowments sold in the 1980s. Where losses come to light, there can be issues in terms of changes in provision – what general insurers might call the reported but not settled kind of risk. An example of that is the LBG loss that started life as a £3.2 billion provision in 2011 but has grown by £13 billion since.

Of course one issue here is that general insurers seek to take on those risks. That is their business, whereas firms do not seek to take on operational risks. They have some control over operational risk exposure. They can invest in controls. But, in general, it can be looked at as a cost of doing business. As a cost of doing business, the operational risk will change as the business model changes. For instance, a lot of life insurers might be launching Wrap products in the asset management space. That brings in a whole new set of operational risks. There are also a lot of changes which can be caused by changes in control frameworks which lead to differences in the operational risk profile. There can be changes down to legislation. That might be just down to changes in actual law, regulations, new precedents such as, for instance, the European Court of Justice rulings on over time and holiday pay which may be costing firms considerable sums.

There is, obviously, technological change. A key one would be cybercrime. Cybercrime has evolved quite rapidly in recent years. Noteworthy is that last year the US health insurance company, Anthem, suffered a loss of 80 million records as a result of a data breach. That is fast becoming an important aspect of operational risk.

You can see that, by definition, operational risk is evolving at a faster pace than the general insurance market can come up with new solutions to mitigate it. This constantly evolving change in the nature of operational risk again makes operational risk modelling so difficult.

In terms of the framework, operational risk models will be useless unless there is a robust framework for identifying and assessing risks. Ideally, the model will feed into that framework. But without that initial assessment of risks, modelling is pointless.

A key aspect identified, in terms of what a framework should have for operational risk models, is the need for a very well-defined operational risk taxonomy. One of the key problems that I have seen in modelling operational risk is that because the loss categorisation is vague, losses can be categorised under the wrong category. It would be the equivalent in general insurance to flood losses being classified under fire; and you have the equivalent going on under operation loss because people do not understand under what category a particular loss falls.

110

Also, in areas such as scenario analysis, a common failing I have seen is where one risk is covered twice over because people do not understand where the risk should be covered. Even worse, certain risks are missed from scenario analysis altogether because people are assuming that they are covered in another category.

I would now like to move on to the data and inputs used for operational risk models. A natural starting point for operational risk models would be internal loss data. Basel II places a lot of emphasis on internal loss data and the loss data approach to operational risk modelling.

However, it is beset by many problems. Key amongst these is the problem of low-frequency/high-impact events that are not in the data. Nearly all firms will have operational risk data sets that will not extend back beyond 2000. Obviously, key elements are missing to operational risk experience.

Again, to use general insurance knowledge, you begin trying to model windstorm losses when you do not have sight of Hurricane Charlie in 1987 because your data only goes back to 2000. There is also a problem that loss is intrinsically retrospective. There is the question of the relevance of that retrospective data in a changing environment.

These are well-understood problems of operational risk loss data. There are other areas we identified in the working party which are worthy of consideration. A key issue we identified is what loss impacts are you collecting, and whether these are relevant to capital modelling. For instance, as part of your loss collection process, you might be collecting details on lost sales or higher lapses. But lapses might already be covered under insurance risk. Lost sales might not reflect on funds. There is a need to understand what losses you should be modelling on your operational risk and what models you should be excluding. I am not saying that you should not collect data on higher lapses or lost future sales. It is just that you need to have the ability to be able to segregate those from your loss data set.

Another issue we have identified is exposure to third parties such as fund managers and outsourcers. On a net basis, the exposure to the third parties will be quite modest because any mistake that they make will generally be indemnified by the fund manager or the outsourcer. But on a gross basis, some of these latent exposures are quite considerable. I think there is a need, as part of your loss data, to track not just your own losses but also losses incurred by third parties in respect of your firm.

Another point on loss data is that when we were going through all the regulatory literature out there, there was some very good guidance given by the Basel Committee on Banking Standards (BCBS) and we reckon that this should be considered by insurers and other firms not subject to the Basel regime. I found it very practical and very useful.

Turning to external loss data, given the gaps in internal loss data, there are obviously attractions to try to expand the data set by using external data. This can either be through confidential loss data sharing schemes such as Operational Riskdata Exchange Association or the ORIC or it could be through providers like Algo First which collates publicly available information on operational risk events.

A couple of issues with external loss data. First of all, I do not think that it fully addresses the problem of low-frequency/high-impact events in data. What you tend to see a lot of the time is rather than large events happening at different times for different companies, all companies behave at the

same time. So, if you look back, in 2010 few banks had London Interbank Offered Rate (LIBOR) fines or payment protection insurance (PPI) mis-selling in the loss data set, but within the space of two or three years, most banks had these kinds of losses come through. It will give you some perspective but I do not think that it is a solution to that particular problem.

Another issue is scaling of loss data to reflect one's own exposure. You can talk about a £16 billion mis-selling loss. That would be completely irrelevant to a small-scale building society. There are challenges as to how to scale losses in an external data set to your own firm's size and exposures.

Unlike internal loss data, you have questions of relevance in the face of changing business models and changing control frameworks with the added complication that it is not just your own business that is changing, it is other people's businesses, and you often do not have sight of changes in the control frameworks of your peers.

Even if these problems prevent you from including external data in modelling, it is still quite useful. External data can help to validate the results of your own internal model. It can help with scenario analysis if you supply scenario analysis experts with external loss events, and it can help in terms of business as usual (BaU) risk management.

A key theme from a review of literature on the topic is that historical loss data on its own is unlikely to be suitable, and needs to be complemented by prospective risks. In general, this should be provided by scenario analysis.

To my mind, scenario analysis is key to operational risk modelling. If it is done properly, it offers the chance to capture the low-frequency/high impact events not in your data, and also can reflect and anticipate changes in your business model and control framework. The problem is that it is intrinsically subjective and it is affected by many types of bias. One type of bias might be an anchoring of scenarios grounded in past losses. That would be quite unfortunate because you would obviously lose out on the prospective view of risks.

Another problem with biases is that people might be afraid of calling out high impact losses, for fear that they may be seen as not properly managing their own risk. It is quite important that any scenario analysis exercise is conducted in an environment without fear of recrimination if any high impacts are identified.

A particular issue you need to be on your guard against is what are called game scenarios, in order to achieve a particular capital target.

To my mind, this is doubly dangerous. For one thing you are tending to pull the wool over the regulator's eyes, which is never a good thing. More to the point, you are fooling yourself about the size and scale of your own operational risk exposure.

In terms of scenario analysis requirements, we have identified a number of particular stages of the scenario analysis process. Preparation is key. There needs to be a lot of detailed preparation going into the scenario analysis exercise.

One of the things that you need to identify is who are the correct experts to consult. Often it may not be the obvious ones. To take an example of processing risk, there will not be just the operations

112

people that might have a perspective on this aspect, but people from HR might have a perspective on payroll risks and pension scheme risks. They should be invited into the process as well.

The subject matter experts involved in the process should be given detailed information on historic loss events, both internal and external, but also details on control self assessments and key risk indicators.

Last but not least is that they should be given detail on the risks to be covered under a particular category, analysing risk by category. As I mentioned before, a key failing in scenario analysis is where people consider risks that they should not cover or they miss risks because they seem to be covered somewhere else. There needs to be a lot of detail in terms of what risks are within the scope of a particular scenario analysis assessment.

Another aspect of the process that we felt worth highlighting is the need for follow-up. Too often, scenario analysis is treated as a one-off exercise. People go into a workshop and produce a loss figure. Unfortunately, a lot of those loss figures are quite crude so there is a need in the scenario analysis process to set aside time to address this shortcoming and follow-up loss estimates and firm up on assumptions made in the scenario analysis.

Given the subjectivity of the scenario analysis, independent review and challenge is key. The Basel Committee suggested this should be done by the third line internal audit. I believe there is a case for second line risk management to participate, particularly if they can demonstrate independence from the process. Having your second line involved in scenario analysis review can help the risk management function understand their own risk profile a lot better.

Documentation has been the bane of actuaries everywhere. It is key to scenario analysis. The key point here is that they should not document just the particular scenario chosen, they should also document those risks and scenarios that were discarded in favour of your chosen scenario because that way you can help evidence-based consideration of all the risks within a particular category.

Finally, a key element would be executive ownership and sign-off. To my mind, an ideal would be where the executive would review and then approve resource. But, also, those individual executive members would take ownership of a particular scenario. For instance, the chief information officer would take ownership of IT failures. That responsibility can help ensure that scenario analysis receives the attention that it deserves.

Once you have modelled the individual risks using loss data and scenario analysis, a key issue then is how you aggregate those risks. A key determinant of your operational risk capital would be correlation assumptions.

We could try to derive the correlations using past data, but the conclusion of the working party was this is unlikely to be adequate given the lack of data. Spurious results could emerge; there could be random error. You might miss tail dependency because of lack of data on high impact events.

When you are dealing with low-frequency events, historical data can systematically underestimate dependencies.

Our view is that you need to supplement empirical data with expert judgement. That is easy to say but quite difficult to do. How do you elicit that expert judgement? It would be impractical, for instance, to

ask subject matter experts to populate a 20×20 correlation matrix with 190 separate assumptions. Aside from the effort required to populate, review and challenge it, to obtain executive approval would be prohibitive. It is also likely that it would not satisfy the criteria to be a valid correlation matrix.

Among the ways around this that the working party identified was to group risks at a higher level. For instance, if we look at the seven higher level Basel I categories, and just determine correlations at that level, you end up with just 21 different correlation assumptions between the seven different Basel I categories.

The problem with that is that you lose out in granularity. There may be low-level correlations in general between broad categories. But individual sub-risks might be very heavily correlated.

Another idea might be, as you go through the assessment of individual scenarios, to try to identify how they might be affected by common causal factors such as high staff turnover or the impact of a flu pandemic. Based on the impact of these causal factors on individual risks, you could infer correlations and dependencies.

Whatever way you do decide to elicit expert judgement, the important point is, as with risk scenario analysis, you need to have an independent review and challenge process, and ideally executive sign-off of the results to ensure expert judgements are reasonable.

Turning to other inputs to the model: A number of times we were making assumptions regarding risk mitigation. There is a need to ensure that these are all backed up by extensive research. You might be allowing for buildings insurance or professional indemnity insurance in your operational risk assessment. It is important to do research on the policies that you have and the extent and term of the cover.

Similarly, with outsourcing agreements you make an assumption that you will be passing operational risk losses on to the outsourcer. That needs to be validated against the legal agreements in place. Also, it needs to be validated against the financial strength of the outsourcer – and whether outsourcers are able to compensate for mistakes they make.

Sometimes when we are modelling operational risk, we conduct it at a higher level than legal entity level. We might model at a country level. Then there is the question of how do we allocate that back to the legal entities within a particular country? That is obviously quite an important assumption. Again, there is a need to do research on this aspect and to look into the legal agreements between, for instance, your service company and the various insurance subsidiaries and banking subsidiaries.

I have seen examples of service agreements between a service company and an insurance company where certain types of loss were excluded or prohibited from being charged back.

For UK life insurers, there is the issue of with-profits funds and the assumption made as to what extent we can charge back losses to with-profits funds. The key thing here is to ensure any assumption is consistent with the Principles and Practices of Financial Management (PPFM). Ideally, we would recommend that these assumptions are cleared with the with-profit actuary and the with-profit committee.

To sum up the findings of the working party, from our review of literature the key finding emerging was that loss data is not good enough and needs to be supplemented by a prospective view and, generally, scenario analysis.

114

Our review of all the regulatory literature has highlighted some very good practice in some of the regulatory guidance, particularly the BCBS guidance on advanced measurement approach (AMA) models. BCBS 196 (which we heartily recommend) ensures matters not covered by Basel are considered.

A prerequisite of any operations model is a sound operational risk management framework, and having a very detailed taxonomy so that there is no ambiguity in terms of how operational risk should be classified.

On internal loss data, there are issues of events not in data; and we would particularly call out whether you have clarity on what loss impacts you are collecting and whether they are relevant to capital modelling. We would also call out the issue of third party losses and any latent exposure that might be building up in third parties.

For external loss data, we know there are problems with low-frequency/high-impact events. We know there are problems with the relevance of external loss data given the different business models and different control frameworks of peers, and also the issue with scaling.

For scenarios, we would stress: the need for a lot of effort being put into preparation; the need to have proper follow-up of scenario exercises; the need for a strong and robust review and challenge of assessments; and the need to have executive ownership and sign-off of scenarios to ensure proper buy-in from the organisation to the scenario analysis process.

For correlations, dependency assumptions or key assumptions, empirical data is unlikely to be adequate and needs to be supplemented by expert judgement. We accept that there are a lot of challenges in eliciting this expert judgement.

Finally, we note that there are a lot of other assumptions in terms of risk mitigation and legal entity allocation which need to be supported by extensive research into things such as policy provisions, legal agreements and, for with-profits consistency, the PPFM.

The Chairman: Thank you, Patrick. The discussion is open to the floor.

A member of the audience: The Basel Committee has just come out with a new consultation paper. The consultation paper has proposed the abandonment of internal models for operational risk in banking, moving to one based on a combination of business volume and the loss experience from the company.

So, given this, I should like to ask you some questions. I have four questions. Given the difficulties that you have drawn attention to with loss modelling, all the data and the correlations, do you think that they have the right approach by going to a standardised model?

Question two is that their model, or their formula, is super linear on business volumes. That is, the bigger your business volume, the greater the percentage of that business volume you have to set up for operational loss.

The third point is that they have a multiplier based on your own operational loss experience. If a company has had a bad experience and learnt from it, it may be less likely to have claims in future, or

perhaps it is just accident-prone and is more likely to have losses. What do you think about that? Do you agree with the idea that it is a sensible approach to take into account loss experience?

Finally, there is no scope for exercise of judgement. You can all think of things like complexity, stability, and staff turnover. We could probably put a short list together of factors that we would think impact on operational loss probability. Do you think that that should be appearing in a formula? If so, who decides at what level the factors should be assessed? Should it be the company, the regulator or anybody else?

Mr Kelliher: That is very interesting. The change was made at the start of this month. I think it was trailed for a while. We had responded to the previous consultation back in 2014. However, it seems that AMA has definitely gone. There does not seem to be any scope for it to be retained. So we are not really consulting on that.

Are the BCBS correct in that? I am not convinced. To my mind, it has developed a very complex standard formula. I am not sure that it will capture all the subtleties of operational risk exposure for different types of institutions. If you think of banking in a Swiss private bank, there might be issues in terms of money-laundering and tax evasion. If you consider that with the UK retail banks, there will be different kinds of operational risk profiles for each. It is a very complex formula. But I am not convinced that it will always give you the right result relative to the risks.

That will result in two things. One is the operational risk capital might be too low for certain institutions, particularly where they do not have much other risk, credit or market risk.

The other problem is where there is too much operational risk, capital, relative to one's profile. My thinking there is there is a tendency in Basel, and regulators in general, to beef up capital requirements in general with systems like the leverage ratio. If we are not careful, if we raise capital requirements disproportionate to risks, we will have broader economic impacts. If you increase the capital requirements for operational risks, that will obviously reduce the amount that can be lent by a bank given the equity base. That will have a wider impact on the economy.

So, basically, I am not convinced about the wisdom of abandoning the AMA. One thing that might have influenced them is that the AMA up until now has been compromised by the fact that the standard formula approach was too low anyway. I think, with the benefit of 20/20 hindsight, we would probably say that operational risk capital should have been three or four times as high as it was. The problem when you have a very low standardised value is that will have compromised AMA models themselves, because if you had an AMA model which produced significantly higher capital values than the standard, your board would just throw it out and say this is nonsense.

I have a feeling that that may compromise not just the AMA model but also compromise scenario analysis by big banks. Big banks could say: "We could lose 3 billion on PPI. Hold on, our standard for capital is only 2 billion". That might have been rejected as well.

There were issues with the AMA, but I think that was probably a function of the standardised accounts being too low.

In terms of the superlinearity point, I think for insurance there have been quite a lot of small life insurers who had near-death experiences during the pensions mis-selling scandal because of the

116

extent of their mis-selling exposure. I am certainly not clear from an insurance point of view whether it is superlinearity.

If we look at things such as pensions mis-selling, it affected quite a lot of very small offices, probably disproportionately when compared to larger offices. I am definitely not convinced about the superlinearity argument for insurance.

The multiplier by own experience of losses: Like you said, sometimes banks can get unlucky, you have one big loss and then you suddenly have this albatross around your neck in terms of the capital requirement. I can see arguments for and against that. Some banks are accident-prone.

Being Irish, I know Allied Irish Bank was a classic case because it had an issue in the 1980s with a rogue trader in its insurance arm in the London market. They followed that up 15–20 years later when they had a rogue trader in its currency market which cost them 600 million. There have been multiple instances where credit risk and a lot of operational risks have come to light as part of its recent rescue by the Irish Government.

So, yes, I do think certain banks are accident-prone.

Finally, in terms of expert judgement, one place you could apply expert judgement is in an AMA model. Unfortunately, that is no longer on the table. I think the AMA gave banks an opportunity. If the standard was too crude for them they had a lever they could factor in.

I understand that they are also not allowing for insurance: things like professional indemnity insurance, buildings insurance and the like. I am definitely not convinced about the wisdom of basically cutting off the support that the London market and other general insurers can provide to banks and not giving any capital credit for that position.

Mr A. J. Couper, F.I.A.: I am also probably not convinced overall, particularly from an insurance standpoint. In some ways it would definitely make life easier. We spend a lot of time in the company I work for trying to model operational risk from the ground up. Just having a standard formula would have made life a lot easier.

When we are thinking about model approval and partial model approval, we have to think carefully about this aspect. There are arguments to say you need to prove, or otherwise, why the standard formula is still appropriate. If you are going to use the standard formula, you need to say why it is appropriate to use it. It is not always easy to say why it is. In some ways, modelling it from the ground up is another way; and arguably a better way and more applicable to the organisation for which you work.

I also think by applying standard formulae you are losing something in terms of the internal control framework that you have been building up in terms of the first line, second line and third line of defence.

Personally, I would not want to take anything away from all that hard work and effort that has gone into building the risk management and governance processes that we have as an organisation.

Mr C. M. Smerald, F.I.A.: I would add that one thing about operational risk is that it is highly complex. There are many things going on, and the nature of complexity would be that one rule is not enough. You need to look at it from different points of view.

Patrick's comment about the standard formula being inadequate previously was interesting. It does feel as if the issue is more with internal models: the ground-up modelling does not necessarily give the right number. It can be manipulated. But the good thing about a model is that it helps you understand your business. It seems to step away from having pressure on people who understand the business is a step in the wrong direction.

It is always good to have a benchmark and figure out how and why you are different. For example, with multipliers for insurance, you can look at data and see if there is a relationship. It is a statistical question, in part.

I am in favour of judgement. I work for AIG. We had a very serious loss. It was not categorised as an operational risk loss somehow. In my opinion it was. There is a data question of how you define losses that could very much alter the position.

Mr C. O'Brien, F.I.A.: I should like to make comments in three areas: the distinction between operational losses and operational risk; the regulatory requirements on modelling; and the need for operational loss and risk analysis for the management of the business.

First, my concern is whether the work to manage operational risk is proving successful. Obviously businesses have operational losses where people and processes do not work perfectly. So, for example, if you are retailer you know that you lose stock through shrinkage, including theft by customers or staff.

Firms know that they cannot eliminate operational losses and we would expect that their financial forecast includes an assumed level of losses in future. Such losses may be regarded as BaU. I realise that the paper is concerned with the impact of operational risk on capital rather than BaU operational risk management.

This might suggest that BaU operational losses do not need to be included in the analysis for capital purposes. On the other hand, there can be benefit in recording them if the firm is concerned to reduce operational losses and understand why they are arising, although clearly there is a cost involved.

There needs to be a judgement about what to include.

Operational risk reflects the potential for operational losses to vary. High risk is undesirable, especially when the prospects of losses are high. It makes it more difficult for a firm to plan and to coordinate its activities and high losses can lead to financial difficulties. Therefore it is the variability of operational losses that is important for operational risk.

However, the major operational losses that arise, such as PPI mis-selling and the LIBOR scandal, are, by their nature, matters that are very difficult to cope with by modelling. I wonder to what extent do banks' operational risk models give them any insight into those major losses.

I read in the paper that the operational risk exchange association has a record of over 450,000 loss events. The magnitude suggests that some BaU losses are included. Is PPI mis-selling recorded once or once per company, or what? Anyway, recording and measuring operational losses is clearly a big industry.

I just wonder to what extent all this effort has led to lower operational losses; and, since we are here concerned about operational risk, whether it has led to a reduction in the variability of operational losses. Any information on that would be welcome.

118

Second, on to regulatory requirements, since clearly these contributed too much of the effort that is now taking place. As has been said, the Basel Committee is now proposing to abandon the AMA for banks' operational risk capital requirements. The committee found that the inherent complexity of the AMA, and the lack of comparability arising from a wide range of modelling practices were leading to a lack of confidence in the outcomes.

To go back to Lord Turner's review of the global financial crisis, where he referred to mis-placed reliance on sophisticated maths and questioned whether there were fixable deficiencies or instead inherent limitations, in operational risk the banks will have simpler calculations in future.

Perhaps there were deficiencies in the complex approach that were not fixable. Perhaps we will see other conclusions of a similar nature in future. No doubt people will question Solvency II in the context of Lord Turner's remarks.

Finally, I reflect on how managing operational losses and risk may be different in the absence of rules that are laid down. For example, the paper notes that loss of new business profits should be excluded from operational risk capital calculations so the capital modelling is carried out consistently. But where a firm is making decisions in managing its business, it clearly must take loss of new business profits into account.

In one case, where a life insurer had reduced the unit prices in the fund where policyholders expected that prices would not reduce, the sales personnel forecast a large drop in sales if policyholders were not compensated. That was an important part of the decision-making process. Clearly, operational loss management in practice must take account of new business profits. This case also included issues of whether general insurers were obliged to compensate the life insurer for its losses. While regulators may have good reason to limit the benefited capital modelling that can be taken for the insurance of operational losses, good management will take account of commercial realities that provide for more benefit than regulators' rules for capital provide.

It does strike me that there is a difficulty here. Regulators want to know that a firm is using its model in practice. But the model that a firm would ideally want to use is not the one that the regulator permits.

If the Basel proposals for abandoning the AMA go ahead, I wonder how banks will respond. Will they run models that are consistent with how the business should be run as opposed to models that regulators wanted them to run? Or will banks scale down their modelling saying that they were more sophisticated than they really needed?

Mr Kelliher: I refer to the issue about variability of losses and whether models can ever address things such as PPI mis-selling and LIBOR.

I think you could, if you have the right framework not just for loss data collection but also scenario analysis. That is why I think scenario analysis is key. To my mind, what probably failed in terms of the provisions banks were making when they were incurring these exposures was that their scenario analysis was not robust enough.

If we take PPI mis-selling, that was a problem that was known from the mid-noughties; the same with LIBOR. The *Wall Street Journal* highlighted this in 2008.

In terms of modelling, I think there are losses out there: things like cybercrime and the risk of the Anthem style data breaches. My view is that we can model these things if we have the right antennae to pick up on the risks as they emerge. In 2005 I think the competition authority started its investigation. The scenario analysis for banks should have been to say that if this goes ahead and we have to deal with mis-selling, what will that cost us?

A lot of the banks continued selling PPI well after they saw the investigation starting. I think that highlights an issue that scenario analysis was not done properly. I do not think the catastrophic risk that was there was called out.

In terms of the focus on capital compared to BAU losses, I think that there is a very important point in the general leakage of losses. I think for banks, certainly with the AMA, there was an exclusion. You could have a capital exclusion in respect of BAU losses to the extent that you could expect them to come through in your budgets. There was a lot of effort in trying to identify what BAU losses were incurring. It is something you probably should be trying to do anyway to try to make it a leaner, more cost-effective business.

Moving on to the regulatory question and the AMA, I have made my feelings clear about the wisdom of this. One thing I would say is although AMA is going, obviously there is a Pillar II regime in the UK by the prudential regulation authority (PRA) which I think requires inputs, scenario analysis and loss data. Even though the context modelling part of the AMA has gone, you still need to do scenario analysis and collect loss data to satisfy the PRA's Pillar II requirements for operational risk.

I suspect with no capital gain from investing in complex models, there may be some winding back on modelling effort which is probably regrettable. The fact the PRA are looking, in the UK context, for scenario analysis and for loss data will keep banks' eyes firmly fixed on the ball in terms of quality of loss data and scenario analysis.

Finally, loss of new business profits excluded from the modelling. I think we should definitely capture things such as lost new business profits and some of these reputation damage aspects as part of our general loss data collection process.

I think the important thing is the need to be able to split them out and exclude them from the model. I wholeheartedly agree that you should collect them to inform your BAU risk management, and also to inform the stress and scenario testing. And as well as operational risk, you have the stress and scenario testing requirement under Pillar II.

Some of these impacts can be quite useful in terms of understanding the impact of a particular reputational damage on future sales. That can be important for the type of modelling that you suggest in terms of on-going viability.

Mr Couper: I think that the framework we have in place needs to try to help support the business in identifying the operational risks that it faces. Recording losses will certainly help. You need a process to identify them.

Clearly, there is a human psychology element to admitting failure through operational loss. We need to keep this aspect in mind in terms of how we might model it, not only in terms of the body of the distribution but also the error in the tail.

120

With regard to the tail, there are so few events that a good framework and a good process will help the company in its ability to think about how exposed it is.

One of the important things as we have seen with the developments of cyber risk is to keep thinking through everything to which we could be exposed, and not only thinking about cyber risk in terms of what it is and how to quantify it, but also thinking about what is the next cyber risk? What else is there out there that we just do not know about? I think having the process in place will help to address that aspect.

Mr Smerald: We had discussion in the working party over many of the issues you raised. Ultimately, we decided to have mostly a capital modelling focus but we tried to be permissive with regards other things. We are extremely sympathetic. It is hard enough to get companies to think about the capital bit, much less getting into business and strategic risk.

Also, as actuaries, we can oversell our models. There is certainly a risk there.

I have a question back to you. Is this something that you think maybe we got wrong or is it more of a question of degree where you think we could have said more in some places?

Mr O'Brien: What I am interested in is the variability of operational losses. If you knew, like a retailer, that 1% of your stock was going to disappear through shrinkage, it is not an operational risk. You know that that will happen.

If you are a financial services firm, if you have things that go wrong all the time, X% say, then that is like a fixed expense. It may well arise in your accounts as an expense every year.

It is the variability of operational losses that seems to me important in determining what capital you need to hold. That is regulations aside. It is the variability of operational losses that seems to me to be important. Clearly, when you have big issues, as the banks have had, like LIBOR and PPI that means that their operational losses are very variable in that respect.

What I have not quite got a feel for is to what extent (you may indeed regard those as outliers) operational losses vary significantly over time, which is why capital is needed to support them. I appreciate it is also potential. I appreciate it is the potential for losses to vary over time, and hence you need to take account of events not in data and unknown unknowns.

But it would be of interest to know, even on the basis of experience to date, to what extent do operation losses vary over time.

Mr Kelliher: In terms of operational risk, I am not sure we really know. We only have about 15/16 years' worth of operational risk data. There have been some very big impact events: PPI mis-selling; mortgage endowment mis-selling; pensions mis-selling; and LIBOR.

In terms of variation over time, maybe, if we had data going back a bit further we might understand a bit more whether we are living in a time of exceptional losses.

We do not really have the data, but there is a lot of anecdotal evidence that some losses are not new. Certainly rogue trading has been going on. You had the Barings's event, and you have had quite

a few rogue trading losses down through the years. There have been various conduct risk losses down through the years as well.

Are we in some new paradigm where everything is going to lead to much higher losses, or has there always been some kind of historical norm?

The Chairman: Patrick, you mentioned relevance in your presentation. I think that is key here as the environment continually changes over time. Therefore, it can be very difficult to say are the losses we have experienced in the past are still relevant going forward and hence the importance of scenario analysis.

I think the point Chris (O'Brien) made around a certain level of loss being baked into the BAU expense base is an important one. While that is there and while that may be projected out, it is still important to be aware of that and to be managing that down. If you can put controls in place which will manage that, then clearly you could potentially create value for the business.

Mr Smerald: My takeaway from your question is we need to really pay attention to the attritional side of operational loss.

My work in claims analytics convinces me that, yes, there is fluctuation. Certainly in terms of claim behaviours we do see some changes with, say, the economic cycle. You will get more surety losses in tough times. You will get more failure losses in tough times.

So there are things that affect people's behaviour; but also there is a growing trend that we are seeing with what I will call risk management technology and "heat".

"Heat" is the degree of ire and anger people might have at your conduct. I think the level of "heat" on insurers and the financial services industry is increasing. By risk management technology, the ability to catch us out for having violated something has increased. So people are angrier and it is easier to catch our mis-doings.

Mr K. M. A. Tawfik, F.I.A.: You mentioned that there are many different approaches to operational risk methodologies. I was wondering if you have seen any convergence of the taxonomy or the distributions chosen, or correlation approaches, within your work?

Mr Kelliher: There was a useful survey that was done by the Institute of Risk Management's internal model industry forum which was quite good in terms of the state of current modelling practices.

One of the key things was, for instance, separate frequencies in severity modelling, certain distributions are quite common, and I get the impression more modelling approaches are involving loss data and scenario analysis.

In terms of dependencies, I think that is a bit tricky. I am not sure we have really landed on that one yet. I have seen a study on dependencies in the paper back in 2008. That was based on their data set. That is the only paper that I have seen.

That paper gives a very good state of the art in terms of operational risk modelling, certainly for individual risks.

122

Something that the working party might turn to in the future is to carry out a bit more research into correlations and dependencies and dependency benefits.

The Chairman: The actuarial profession has sponsored research around complex systems theory and Bayesian networks, and its application to operational risk. Have you seen much take up and adoption of that among insurers?

Mr Kelliher: I have only come across one company using Bayesian networks. However, Neil Cantle, who is head of the risk management thought leadership subcommittee, is very much at the forefront of these techniques. In terms of using Bayesian networks for operational risk modelling, I am aware of one company which has definitely applied this technique. There are probably a few others. That could be an interesting topic to research.

The whole point of Bayesian networks is you obtain a distribution of the entirety of losses so it could be useful for understanding the aggregate distribution of operational risk losses. That is useful to keep in mind.

Neil produced a paper on this subject back in 2010/11. His paper is on the actuarial profession's website and he has also produced his own paper for Milliman.

Mr Smerald: Certainly capital models are becoming more structural. Instead of having correlations, you have inputs; variables; direct wired impacts of factors etc. I think that trend will continue.

An interesting extension of that are scenarios. Scenarios are this static thing that you work up and the question becomes: if the conditions change, is your scenario still valid or not? Some interesting thinking around how to make your scenarios reactive to changing circumstances is required. That is a baby step towards structural modelling.

Mr R. Kelsey, F.F.A.: I just want to ask whether you think a one year time horizon of capital modelling is detrimental. I am thinking that most of the large operational risks have been late on exposure, and including some of those which are not operational. For example, we have asbestosis, PPI and VW's emissions scandal. These have built up over a large number of years – even Tesco's advanced billing. Suddenly, over one year you have a monstrous liability which you cannot do anything about because it has built up over a number of years, and you do not know about it.

Mr Couper: I think that the difficulty there is the distinction between underwriting or insurance risk and operational risk. Some of those are perhaps more underwriting risks for an insurer rather than operational ones. We deliberated on that as a working party. The distinction between what is an underwriting risk and what is an operational risk is a difficult one to analyse.

Mr Smerald: It is my personal view that everything has operational risk in it and you do not know until much later. So my advice would be to treat everything as an operational risk and then figure out what portion it represents.

Mr Couper: I will ask a question around the place of models. Is it really around operational risk quantification? Is that a good place to use them?

Often that does not incentivise the right behaviours around the management of operational risk, so how best to manage operational risk?

Mr Kelliher: I think for me the value of models is not so much the end result and the mathematical sausage machine, but what goes into it in terms of loss data capture and scenario analysis.

I suspect what happens sometimes with models is people are looking for a particular result. They backfill and then you end up with this horrible situation where people are more or less discarding scenarios because they do not fit with certain preconceived results.

Done rightly, a good modelling process can add understanding. If your intentions are not pure, they can create a lot of fog in the way of understanding the risks.

Mr Smerald: One of the points mentioned was one year and how does that make sense. The answer is that you do not have to apply one year. If you think something has structurally changed in the environment then it is just the world has changed and you fully reflect it. One year only applies to certain types of instances. It does not have to apply universally.

Mr Couper: I think in terms of how best to manage operational risk, there are probably two aspects. The first one is in terms of the business and what is going on in the business. It is trying to embed a control process. It is arguably trying to minimise (not sure it is eliminate) and control the risks as best we can.

Then in terms of how we are going to model that in setting capital requirements, you need to take into account the control process. In a sense you almost want to keep the modelling away from the business so that there is no gaming of the model. The business is not totally sure of how the parameters of the model are being set and cannot work their magic in terms of persuading you that the risks are really a lot lower than you might otherwise perceive.

Mr Smerald: And even if you understand the question better, you have benefited from the model, and maybe next time, now you have understood it better, you will do a better job. That is for the purists in the insurance world. I am sure that there are companies like that somewhere.

The Chairman: Patrick, you mentioned obtaining a second line review which seems to suggest you think maybe the first line should be owning the operational risk capital assessment. I have seen some second lines taking the ownership for that. Where do you think it should sit?

Mr Kelliher: I think a lot of time operational risk ends up with the risk management function as opposed to the actuarial function, partly because actuaries do not really want to get involved. A lot of actuaries look at operational risks and say: "We can model that. We can fit various distributions to FTSE data. We can model insurance risk". They look at operations and they say that this is too much.

The actuarial function should take ownership of the operations models and you have risk management providing the review and challenge. I think that is probably reasonable. I think the risk management function has a lot of insight in terms of the actual risks on the ground to say whether this risk is reasonable.

Whichever way you cut it, one side needs to challenge the risk management and vice versa. There needs to be that degree of challenge built into the modelling process. Just as individual scenario analysis, in terms of recording, will be done by the first line, there is a need to review and challenge

124

those things, and when you do the modelling there is need for that to be challenged, leaning towards risk management to do the challenging. The key thing is to ensure there is another function, facing whoever is doing the modelling, to ensure that this is a robust discussion, whatever the results.

The problem with operational risk is you can come up with some very strange figures if you are not careful.

Mr Couper: That sort of responsibility needs to sit within both the first and the second line. You have three elements. You have those who own the operational risks, which should be an executive sitting in the first line. For each of those risks you should have a number of controls in place. Each of those controls will have a control owner. There needs to be a very strong relationship between the risk owner and the control owner.

Then you have risk management. Operational risk management sits there overseeing the process that sits between the risk owner and the control owner and ensuring it is working as effectively as it could. What you have to have is a fluid movement and relationships around that triangle. I have certainly seen from personal experience in the company for which I work, you have a good relationship between operational risk and the risk owners in the business and between operational risk and the control owners. Sometimes the link between the risk owners and the control owners breaks down. That is the bit where operational risk management then plays a part in trying to encourage and make sure that the three elements of the triangle are all working properly.

Mr Smerald: I am no longer directly involved in operational risk modelling at my firm; but I will venture a guess at what I would expect. Except for losses, the second line is doing the capital modelling. There is a very strong model risk management department probably acting almost like a third line. There would be a great deal of review and challenge of the underlying model assumptions and the implementation.

The Chairman (closing the discussion): I think that it is fair to say that operational risk is not one of those risks that boards and insurance companies actively go out and seek. As Patrick said, it is an inescapable risk to which we become exposed through seeking other risks for which management do have a preference. So, given that we are going to be exposed to operational risk, there is then a question of how do we measure and manage that risk.

We have spent most of the time talking about the measurement side. We have talked to the variability of data sources that can be used for that, but also the variability of methodologies used across the industry.

We have talked to the need to combine various sources of data, internal loss data, with external loss data, and perhaps to use those to inform our scenario analyses and blue sky thinking through workshops and other means.

We have also discussed a range of issues with those such as different business models that different firms are using, the relevance over time as the operating environment changes, and our ability to scale all those things and adjust the data over time and from one business to another.

We have talked about the importance of having two pairs of eyes to look at the operational risk capital assessment, and a variety of ways of executing assessment. It is critical to have somebody challenging those people who are responsible for the model.

We have had a lot of challenge. I sense some scepticism, given the scarcity of data and the variability of the modelling approaches around the modelling of operational risk. I think that it is fair to say that the members of the working party here all think that it is a worthwhile exercise to continue doing that; and perhaps some of the other approaches that the actuarial profession have been actively sponsoring in terms of research around complex systems theory could be a way forward.

However, I come back to the management point that a capital assessment in respect of operational risk may not be the right tool for encouraging the right behaviour around the management of operational risk. Often it is risk management rather than capital management about which we need to be thinking. I come back to the example that Chris (O'Brien) mentioned of operational risks that are baked into BAU results and not being ignorant of that fact, but chasing down those issues and trying to manage them out.

With that, I should be very grateful if you could join me in thanking our speakers in the usual way.

126