

SESSIONAL MEETING DISCUSSION

Validating operational risk models

[Institute and Faculty of Actuaries, Sessional Webinar, Tuesday 10 October 2023]

Moderator (Mr M. H. D. Kemp, F.I.A.): We are here for a discussion of a paper on validating operational risk models. We have the undoubted expert in this field as far as the actuarial profession is concerned, at least in my view. It is Patrick Kelliher, who is going to present on the topic. But what I am also hoping for is that you will have questions if you are in the room, or you can highlight and add questions to the Q&A if you are online. My name is Malcolm Kemp. I am a member of the Risk Management Board of the Institute and Faculty of Actuaries (IFoA) and this is one of the sessional meetings that we are holding on behalf of that board.

Without further ado, I should introduce our speaker. Patrick (Kelliher) is a Fellow of the IFoA with over 30 years of experience in financial services, predominantly in the life insurance sector. But he has also been involved in a variety of other fields. He is also on the Council and Management Board of the IFoA. He is a Chartered Enterprise Risk Actuary and belongs to a number of working parties within the actual profession, including the Operational Risk Working Party, and it is that working party through which this paper has been produced. We are privileged to have Patrick here and to hear what he is going to say about operational risk modelling and how to validate the models.

Mr P. O. J. Kelliher, F.I.A.: Thank you. Good evening, everybody. I am here to talk to you about the paper on validating operational risk models with specific focus on the three most common approaches that you tend to find in operation risk modelling; namely, the loss distribution approach (LDA), based on historical loss data; the scenario-based approach (SBA) based on scenario analysis; and also touching on causal factor based models like Bayesian Networks (BN), which are growing in importance. This paper is the third in a series of papers produced by the Operational Risk Working Party. It builds on previous papers in terms of inputs to operational risk models and aggregation and dependency for operational risks. This also reflects a lot of the practical experience the members have had in calibrating and validating operation risk models. In terms of the talk tonight, first I am going to set the scene and talk about the challenges of modelling operational risk before saying how an actuary or any other risk professional might want to validate this.

In terms of the challenges, the first thing to note about operational risk is that it is a very diverse category. It covers everything from processing errors to cyber-attacks to money laundering failures to employee relations losses. The Basel framework, for instance, has seven Level 1 and twenty Level 2 operational risk categories, but you can often identify more than one hundred Level 3 sub-risks underneath these.

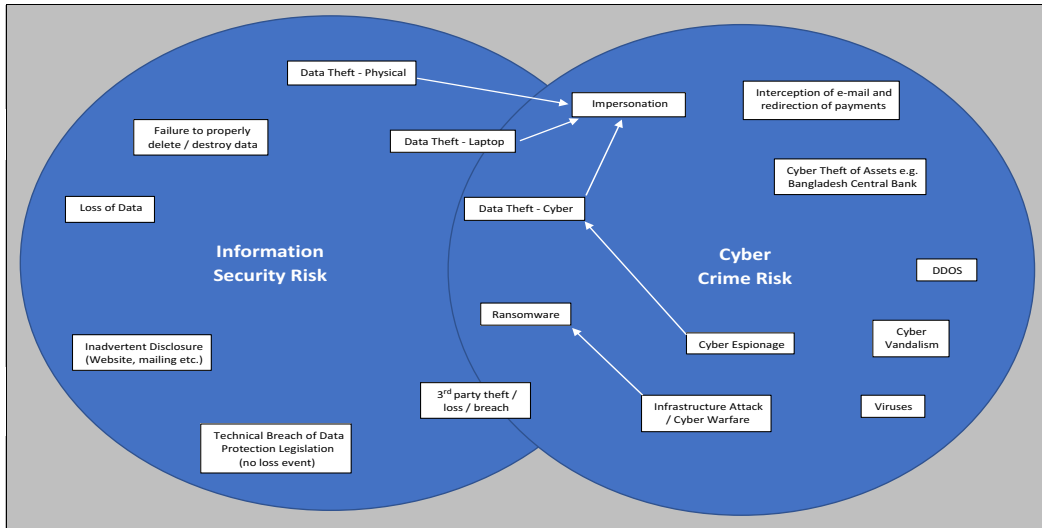


Figure 1. Challenges of modelling operational risk

Figure 1 is based on two generic Level 2 categories, one for Information Security and the other for Cyber Crime. I can identify about eighteen or so sub-risks underneath those two categories, each with their own particular loss characteristics. Another problem we find is that there is considerable overlap between categories. One challenge with operational risk is to understand the demarcation line. What you are modelling, and under what category, can be quite an important problem to solve.

You can identify potentially hundreds of sub-risks. It is usually impractical to model every single one, and so trade-offs are required. With LDA, for instance, we might group all of our losses for different sub-risks under each Level 2 category. But then we introduce heterogeneity into the data because you are grouping losses from different risk types. For SBA, what tends to happen is that you carry out scenario analysis at Level 2, but then the problem you face is that you only focus on one or two sub-risks; hence the question arises as to whether we are properly covering all the other sub-risks.

The other problem within LDA and SBA is that, having modelled individual categories of risk, we then need to aggregate them in some way. Another challenge around operational risk is the very bespoke nature of it; operational risk profiles vary from firm to firm. We have differences in profiles between banks and life insurers, but even within life insurers operational risk profile might vary widely between an annuity specialist and a unit-linked office. There are different business models as well, for example, different distribution channels. You might have different outsourcing models and that will impact your operational risk profile, as will governance arrangements and control frameworks. Unlike equity risk, for example, operational risk is very idiosyncratic.

Another challenge is the evolving nature of many operational risks, particularly cyber risk in recent years. Ransomware attacks, for instance, have grown exponentially, which is a challenge to model. But even for risks which are not changing in general, firms' exposure can change quite radically because of decisions made. For instance, the decision to enter the annuity market will introduce a whole set of operational risks that you may not have encountered before.

There is also the problem of legacy exposures. For example, mortgage endowment claims in the 2000's arose from mortgage endowment sales made around 15 years previously. As with general insurance, we have a huge problem associated with a long tail, i.e. a long time for claims to emerge.

Another problem is insufficient data. Often, life companies might only be collecting data for 20 years or so, and this might not be enough time, particularly for very low-frequency, high-impact

losses. External data can help. But even then, there are issues in terms of its relevance and scaling. And because we have poor data, there is a higher reliance on subjectivity in the modelling, which is a challenge not least because there is a potential for bias. This is particularly problematic when you are trying to model regulatory or economic capital, and there might be pressure to produce a low figure. This is something the validator needs to be conscious of.

What might a validator look at when validating an operational risk model? They should start by trying to understand the actual profile of the firm in question. No two firms are the same. One firm might have a particular exposure that is not shared by any other firm in that industry. There is no “one-size-fits-all” approach to modelling and validation.

Something else for validators to consider is changes in risk profile. What I have often seen is that the model has not been updated to reflect changes in the firm’s profile. They also need to be conscious of developing areas like cyber risk and potential legacy exposures. For example, are there any mis-selling exposures in the woodwork?

The other thing I would always look for as a validator is an articulation of why they have chosen this model. Why have they chosen LDA or SBA, and just as importantly, why have they rejected other approaches? Validators should be aware of the model governance framework in which the operational risk modelling was carried out. If it is not very robust, then the likelihood is that substandard practices may be present in the operational risk model. Expert judgements should be tracked.

Shoddy documentation is usually a marker of shoddy development practices and is a red flag for me as a validator. If you are validating a Solvency II Internal Model, the fact that you, as an independent knowledgeable third party, cannot understand the model means that there has probably been a breach of Internal Model rules regarding documentation.

Model uses are another area around which I would want to take comfort. Are the inputs and outputs being used in the wider business, or is this an ivory tower exercise undertaken by actuaries? Sometimes there are challenges. For example, sometimes model outputs are not very sensitive to changes in controls because of granularity issues. When it comes to operational risk, often the journey is as important as the end result. The inputs like the actual loss data and scenarios have a lot to offer to the business. I would always look to see if that information is being used. Or is this just something that people are doing to humour the actuaries?

Finally, culture. This is very difficult to measure or validate, but obviously something to be aware of. For instance, is there a very aggressive culture in terms of modelling and trying to always get the lowest figures?

I would like to talk about some general modelling considerations. The first one is kurtosis and size of tails. Operational risks are generally quite fat-tailed with potential for severe, if not catastrophic, losses. But I have seen models that are very thin-tailed. That could arise, for instance, from LDA, where you do not have large loss events in your data. For SBA, a particular issue is where loss estimates are quite close to each other. You can end up with a very thin-tailed distribution, so that is always a good place to start looking.

Simulation errors are good to look into as well. In a lot of operational risk models, you model severity and frequency separately and then use simulation to model the combined distribution. Because of the low frequencies, you often need quite a lot of simulations. One million simulations of loss are common, but even that may not be enough for certain low-frequency, highly fat-tailed risks.

Caps on modelled losses are quite frequent in my experience. I do not see any particular issue with capping the severity distribution provided there is a rationale for it – i.e. it is grounded in boundary constraints. If you only have a limited number of policies, there is only so much money that those policies can lose, for instance. However, any caps on losses need to be justified in terms of boundary constraints.

Regarding recoveries, like insurance recoveries or recoveries under outsourcing arrangements, best practice should be to model losses and recoveries separately. An issue with modelling net

losses is that you can implicitly model recoveries that are greater than the amount covered under an actual insurance policy. For models of recoveries, we should look at, for example, the sum insured, to see the maximum that can be claimed on their policy, or the maximum indemnity under outsourcing agreements, which are caps on how much you can recover. Also, we need to model the probability that you may not be able to make a recovery. For instance, if you have looked at modelling damage to physical assets, it may be the case that the damage is caused by a peril that is not insured.

One issue relating to life insurance is where you charge losses back to policyholders. That needs to be justified in terms of the Principles and Practices of Financial Management (PPFM).

Another general issue is recurring losses. Sometimes the operational risk model focuses on very low-frequency, high-impact events, but then you might have very high-frequency but low-impact losses recurring quite frequently that you can expect to experience loss on nearly every year. Sometimes they are covered under maintenance expense assumptions, but not always. You cannot assume they are covered by maintenance expense assumptions. You may need to look into where those losses are being covered in existing modelling.

Finally, there is some good benchmarking out there, particularly in methodology, produced by ORIC, KPMG, EY and others. But one thing to bear in mind is that benchmarking gives less of an assurance than it does for, say, equity risk. The fact your operational risk requirements appear similar to peers may not mean much if your operational risk profile is markedly different from peers.

A few comments on specific models, taking the LDA approach first.

The validator should look at how it is covering risks; one approach is to see how the actual loss data used maps to various Level 2 categories, to see how much data there is for each Level 2 category. Are there gaps in the data? How does the model address those gaps to make sure that there is proper allowance for some sort of loss arising under a particular loss category?

There is also the case of Events Not in Data (ENID). Even with 20 years' data, you may not have certain very low-frequency, high-impact events in the data. You need to understand the potential for large losses to arise in that category, and whether that has been properly allowed for.

You need to consider the relevance of data to current risk profile. For example, cyber risk is evolving so rapidly that historical loss data might not be much use in terms of modelling future exposure.

There are also other changes in risk profile. We have mentioned that if you enter a new market or a new outsourcing arrangement, that will then create risks that are not in your historical data. Also, there is potential for historical loss data to be no longer relevant. For instance, you may have lots of mis-selling data but you may no longer have any mis-selling exposure. That is something to consider.

Loss components are another issue. One problem I have found is that when you look at loss data, there are certain components to loss. For instance, certain loss figures might include an estimate of lost sales. That may not be relevant if you are considering an economic capital assessment, where you are not making any allowance for new business value. Conversely, the loss data may not capture things that impact value-in-force. You could have a loss event that involves reduction in future charges but is not reflected in loss data.

External data can cover a lot of gaps in internal loss data, but there are problems with that. First is the question of relevance. For instance, a lot of life insurance loss data will include frequent unit processing errors. If you are just writing annuities, that may not be relevant to you. There is also a need to scale. Monetary loss suffered by a large insurer might not be very relevant to a small friendly society.

There is also a question of quality assurance and governance around the loss of data inputs, and making sure that they are appropriately reviewed and checked to ensure that they are accurate as far as possible.

Most of the focus here has been around the inputs in terms of the modelling. It is similar to what you would do with equities or bond spreads in terms of fitting distributions. But there are

two things worth bearing in mind. The first one is a granularity point – to what extent are you grouping different sub-risks together into a category in order to fit a distribution? In doing so, you introduce heterogeneity in that distribution. The other thing to be aware of in terms of sensitivity analysis is that, given the lack of data, sometimes results can be very sensitive to the addition of an extra year or, for instance, the removal of a large data point.

Turning now to SBA, the key challenge I have come across is how to justify that the model covers all the risks in the category? As I mentioned, you could have twenty sub-risks in a category. It is probably impractical to do scenario quantification for all of those. What we end up having is a couple of representative loss points for that entire category. In terms of how I would address the challenge, I would first look at documentation. Is there evidence that the full range of risks in that particular category has been considered? Under cyber-attack, you might focus on ransomware, but there are other types of cyber-attacks. Have they been considered? Sometimes what I like in models is where the representative losses are from different sub-risks. For instance, the typical loss might be a Distributed Denial of Service attack, and its impact on the firm. A severe case could be, for instance, a ransomware attack. I would always look at the frequency parameter. I like that to reflect the probability of a material loss event arising for all the sub-risks in that category, rather than just the particular representative scenarios chosen.

Last but not least, to come back to the issue around evolving profile, we should check to see that the scenarios are reviewed at least yearly to make sure they pick up on any changes.

Apart from the risk coverage, as a validator, I would look at the SBA process. How do we come up with a scenario analysis? What subject matter experts (SMEs) are involved? Are they senior people from across the organisation? Or are they a bunch of junior staff from one particular area? That is obviously going to affect the quality of the discussions around scenario analysis and the outputs.

I like to see proof of really good supporting evidence being given to them so they can make a proper assessment of exposure. First of all, it should be very clear to them what sorts of risks they need to consider for a particular category. It should also ensure that they are fully briefed on historical losses, the current state of controls and future plans that may affect risk profile.

There needs to be a robust discussion on scenarios that could arise. Documentation is key. I like to see robust documentation giving the gist of the discussions involved. What risks were considered? What risks were discarded? Is there any evidence of bias? Is there a bias towards current events at the expense of legacy, for instance? In the absence of good documentation, it is very difficult to form a view on these issues.

Governance is key, because of the subjectivity of scenario analysis. It is important that you have robust governance. You should have second- and third-line review and challenge of results. Ideally, you would also have senior management involvement at this stage, attesting that the assessment was robust.

A particular problem I have encountered is loss estimation. Often you find in workshops that figures are picked out of the air and no further analysis is done on them. The validator should look at these figures and make sure that they did the research in terms of historical losses to see if it is a reasonable loss estimate.

Having come up with your scenario inputs, in terms of modelling, there are a couple of other issues. For severity distribution, you have a typical loss that is generally assumed to be the median of the distribution, and a more severe loss, which may be assumed to be the 90th percentile of the severity distribution. The actual results will be very sensitive to that percentile assumption. If it is 95% or 80%, it will make a huge difference to the results. Also, sometimes there can be a lot of confusion in terms of what the 90th percentile means. If SMEs are not clearly instructed, a lot of times they consider a 1-in-10 event when arriving at the severe case loss as opposed to 1 in every 10 losses, so it is important to ensure their instructions are clear on that point.

I mentioned modelling before in terms of kurtosis, a key point being that the results will be sensitive to not just the size of the loss estimates, but also to the relative ratio between them. Sometimes you can get very large loss estimates, but because they are very close together, you end

up with a very thin-tailed distribution or alternatively, you could end up with implausibly fat-tailed distributions.

Another issue is back-testing of results against historical loss data, which is another important matter for a validator to consider.

Once you model LDA and SBA, you need to aggregate up. Invariably, you need some form of correlation assumptions. There is usually not enough data to derive correlations. Even if you are able to derive empirical correlations, for low-frequency risks empirical correlations may systematically understate the underlying correlation between the two. I suspect there will be some reliance on subjective expert judgements, but the validator should check to see if there has been a review and challenge of these assumptions.

Another area to consider is how certain common causal factors, such as the impact of economic recession or a pandemic, would affect different risks and whether that is reflected in the correlation assumptions. The paper gives an example of some of these factors, and you should also have Own Risk and Solvency Assessment (ORSA) scenario testing. Sometimes in an ORSA, you will see that the scenario has implications for operational risks, and the validator should see if that kind of relationship is consistent with correlation assumptions.

In terms of modelling, you have got to look for sensitivity analysis of different approaches, but I would caution against simply saying, "Oh well, using a Gaussian copula, that is not good enough." Often you find a lot of refinement in operational risk modelling is spurious, especially when you consider the subjectivity of the underlying inputs.

Sometimes, once you aggregate losses, you have to allocate them to particular risk categories or to legal entities. One thing to consider is service level agreements and other contractual arrangements between companies. Sometimes it is not possible to charge back certain types of losses, for instance, employee relations losses between subsidiaries. Also, when considering charging losses to With Profit Funds, these should be consistent with PPFM and ideally the With Profit Committee would review and challenge them.

Touching on Bayesian Network (BN) models, I have not been involved in validating them, but the working party have had a lot of input from Neil Cattle, who has a lot of experience of them. So, thanks to him for this.

In terms of the areas to consider, begin with the causal factors that the BN model is built on, like errors in certain processes, flawed recruitment and the like. Check to see if any underlying cause or factors are obviously missing.

Again, look at the question of risk coverage. The BN model might not be able to cover every single sub-risk but are we seeing risks/losses arising from the model from each of the high-level categories?

A major consideration is sensitivity testing. Sometimes different nodes of the BN model can be very sensitive to a very small change in an assumption; for instance, how flaws in employee recruitment could lead to fraud losses. It is important to "kick the tyres" on these assumptions in terms of data source reliability and validation, or, where using expert judgement-driven assumptions, that these are based on a sound expert judgement framework.

Back-testing is quite useful against historical losses. And, ideally, you would have a light touch scenario analysis exercise running in parallel with your BN models, just to act as a check to see if there are any particular scenarios that the model is not capturing.

The final point to note is IT systems. Some IT packages are better than others. Some packages struggle to deal with very complex operational relationships, which results in a lot of restrictions on what the model can do. Solvency II regulations address IT infrastructure and how that impacts your model results, and this is one example of where such infrastructure is important.

To summarise, operational risk is a very complex area. It is very diverse and there are a lot of idiosyncratic elements. There is not a huge amount of data and so there is a lot of subjectivity. So, it is difficult to model and validate.

To any validator, I would say, “Start from understanding what the firm’s risk profile is, and in particular any changes in this, which often get missed.” Consider also the qualitative aspects as well as the quantitative ones; as actuaries, we like to look at the actual distributions, but one also needs to consider qualitative aspects such as quality of documentation, the model governance framework and the firm’s culture. On the technical side, common issues would be the tail of the loss distributions and simulation error.

Modelling of caps on losses is another area to consider, along with modelling of insurance and recoveries, and recurring operational risks, i.e. recurring high-frequency low-impact operational losses. For LDA and SBA, you also have the subjectivity of correlation assumptions, which need to be reviewed.

That is a whistle-stop tour of the paper. I will now go back to Malcolm (Kemp) to discuss any questions you may have.

Moderator: I would pose a question to the people here or those online. If you have been involved in validation of these types of models, is there anything that you would also want to add to Patrick (Kelliher)’s excellent precis, because one of the key things about developing good models is to draw on a diversity of input. If you have got some input, some insight that you would also want to share, I am in the process of updating a paper that the Actuarial Association of Europe has prepared on operational risk, and we have arrived at the topic of risk culture. So, if anybody has any bright ideas as to how to handle risk culture, to put that into the melting pot of validation, then I personally would be very happy to hear those. Does anybody have any questions at this moment?

Questioner: My question is on a practical basis relating to the Financial Services Commission in Jamaica, which regulates the non-banking financial sector. A month ago, there was a serious cyber breach where 75% of their servers were encrypted. The Chief Risk Officer is a friend of mine and I have been chatting in confidence with him about different things. I proposed to him that, even though they had all the experts coming in to advise them and spending huge amounts of money to solve this problem, they should have an independent third-party expert providing oversight. The oversight must not be connected to, say, one of the commissioners. I suggested a consultancy or someone who does not have any direct connection. I would like to hear what Patrick (Kelliher)’s views are on having an independent oversight.

Moderator: Yes, can I put in a plug? The previous sessional meeting that I came to was on cyber risk, so, if you are interested in cyber risk do check out that material as well on the IFOA website. But over to you (Patrick).

Mr Kelliher: I think there are two aspects to consider. There is the modelling and model validation, and there is the actual risk. Cyber risk is something where I think all organisations are struggling. They can always do with some external help. Specifically, for cyber risk, penetration testing is key. I have been in organisations where an ethical hacker came in and stole thousands of records just when he was in the lobby waiting to speak to us. It was very illuminating. We all thought we had reasonable controls, but we didn’t. So, for both aspects, it is always good to get that third-party view. The other thing I will say is it is always better to look at external data and learn from what others have done wrong or suffered from before you encounter the same challenges. So, that is another aspect. It is getting that third-party view of how good your controls really are and also looking at what has happened to others because you might be next.

Moderator: An online question: is it appropriate to set a minimum threshold on losses? So, I guess, if you were an organisation and you had a shoplifting risk, you might not go to the same trouble monitoring individual shoplifting events if they were only going to be a few pounds. Although presumably, you would want somehow or other to capture that risk if you thought it was quite likely?

Mr Kelliher: Yes. I think this is something we covered in our previous paper on inputs. I am going on memory a bit. In life company losses, often what you will see is lots of small operational losses around processing errors. A lot of the time what happens is there is no huge loss and no huge compensation, but you might give, for example, an ex-gratia voucher for £50. Certainly, you would not be looking to capture those small amounts, but you do need to check and keep an eye on them. It is important to see where they are being picked up. I have seen an instance with one company where there were these small ex-gratia losses and there was an assumption that they would have been picked up in the expense analysis, but when we looked through the expense analysis, they had been deliberately excluded. So, this area of recurring losses, sometimes very small, is important. It makes sense to have certain thresholds, for instance £10,000 is quite common in banks, to avoid being swamped with low-value losses, but you need to make sure that these are not being ignored and they're covered somewhere.

Moderator: Another question: Do you have an insight into alternative ways of dealing with operational risk other than in the paper, perhaps ones that might be more asset management focused?

Mr Kelliher: When it comes to asset management, because they were covered previously under the Basel regimes, the Advanced Measurement Approach (AMA) to operational risk capital modelling might be quite common; and the Basel AMA was more an LDA-based approach. However, asset managers can use scenario analysis just as much as insurers or banks. There are also hybrid models: I have seen one hybrid model where you would use LDA for certain very high-frequency losses and then use SBA for low-frequency losses. Apologies if I cannot share any more insight.

Moderator: You mentioned that modelling operational risk can be a case of spurious accuracy – at what point do you think that might kick in? Is loading a good approach?

Mr Kelliher: Yes, it's interesting. An example of that is if you look at life insurance in Solvency II, a lot of people might say, "Well, I'll just use the Standard Formula" which is just a loading to capital. The problem with that is you need to demonstrate that it is appropriate for your risk profile. You need to do some assessment of your risk profile, not just under expected conditions, but also under extreme events; and carry out modelling of that, perhaps based on scenario exercises.

Moderator: In workshops you might be looking at a 1-in-20 or a 1-in-100 event and somehow you have to scale that up to a more extreme event, or maybe it is part of the LDA modelling that you're doing, but do you have any advice on how to tackle scaling up?

Mr Kelliher: The problem I have encountered has been the reverse, where a typical loss might be a 1-in-10 event. You might have a severe case loss where it is assumed that nine times out of every ten the loss won't be as bad as X. The problem I have found in workshops is people thinking of a 1-in-10 severe case loss as a 1-in-10-year event. So, you can get a lot of confusion. The instructions being given in scenario analysis workshops are crucial. In terms of the 1-in-20 or 1-in-100 event, I am trying to understand what those would entail. Obviously, the trite answer is you could use Excel and say X is the 95th percentile and Y is the 99th percentile, and extrapolate that way, but it is probably a bit more sophisticated than that.

Moderator: If there are questions that you would like more information on, I am sure Patrick is available through social media, if a question has not met your particular need.

A question about management actions. How would you take that into account, or are there any pitfalls in thinking through the management actions issue?

Mr Kelliher: With LDA, you are modelling losses which reflect historical management actions, but the question is are the historical management actions still relevant? If you have strengthened your controls, are the historical management actions still relevant? Within scenario analysis, which is what I am more familiar with, typically when looking at loss estimates there will be a consideration of what actions can be taken. Often, the typical case might assume certain management actions, but the severe case might assume that management actions and controls fail. I think that will be implicit in the actual scenario.

Moderator: I guess that would also apply to other recoveries that you have described?

Mr Kelliher: My preference for recoveries would be to try to model them separately. If, for example, you are modelling recoveries under an insurance or reinsurance policy, there might be a change in the cover, or a decision to lapse cover altogether, or it might not be possible to secure cover. I accept that there is an element of one rule for one, one rule for the other. My preference in scenarios is to allow for management actions in the loss estimates. If you have a cyber-attack, you could put certain things in place to mitigate losses for individuals. I prefer to allow for those, but for recoveries or insurance programmes and outsourcing, where you are reliant on the third party paying up money, I prefer to keep those separate.

Questioner: Would you expect to see the risk register incorporated as part of the modelling approach, regardless of LDA/SBA?

Mr Kelliher: Yes. We talked about the validator understanding the risk profile. I would expect to see, as part of that process, the validator reviewing the risk register. You have a list of all of your risks and controls and what the state of those controls are. I would see that as a key part of how to understand risk profile. You might have some organisations who have particular issues in one area, for instance, employee relations. A lot of the time, it is not a huge feature in loss data, but it would come through on the organisation's risk register.

Moderator: We have got another online question referring to a ransomware attack that resulted in the regulator imposing a capital add-on, instead of a fine. What is your view in terms of applying some kind of overlay? Or how would you deal with that type of regulator response to operational failings?

Mr Kelliher: It's a difficult one. If you are talking about, let's say, an add-on to the Solvency II Standard Formula, maybe in the UK if you had a huge loss and the regulator said, "Well your Standard Formula capital is clearly not appropriate." I think yes, you have to have the add-on, but in the background I would still carry on with my assessment of what the risks are, regardless of the add-on, and see if the Standard Formula plus the add-on is still a reasonable reflection of the risk. I suspect we could be able to say it is probably an overkill, in that the Standard Formula plus the overlay is probably much greater than our risk.

Moderator: Another online question is, if you are allowing for some operational risk losses, such as run-of-the-mill maintenance expenses, how might you validate them? Maybe you have been primarily validating the 1-in-200 event and this is the 50/50 one.

Mr Kelliher: This is again thinking about recurring losses. I think it is about challenging assumptions. The question, for people in the room, is where some level of operational loss is assumed to be captured within maintenance expenses, how would you validate this assumption? As I mentioned earlier, I was in the situation where a senior actuary said "There's no need to validate this, all operational losses are included in our maintenance expense assumptions." It seemed reasonable, but then I looked through the expense analysis and that was the example where there were actually excluding recurring very small losses. We need to test assumptions. If

you are assuming it is in the maintenance expenses, look at the expense analysis and confirm it is in there.

Questioner: I was wondering, how does a firm decide it should use LDA over SBA in modelling operational risk, or vice versa? Do you have a personal preference between LDA and SBA?

Mr Kelliher: In terms of LDA versus SBA, there are two particular aspects. For banks, if they were using the AMA internal model, which is the equivalent of the internal model approach in Solvency II, LDA was required. You did some scenario analysis, but it was pretty much grounded on LDA. Sometimes there is a regulatory driver to a certain approach. The key thing is that, if we are going to use an LDA approach, do we have the data? I must admit, I have tended to encounter SBA due to the lack of data, particularly for the tail events, for those very infrequent high-impact losses. As data gets better, and we get better at integrating external loss data, LDA would be a lot more viable, but even then I would say there are certain issues. LDA is basically historical losses. For certain risks like cyber, which are fast evolving, even if you had lots of data, it may no longer be relevant. Most people end up with SBA or some hybrid just because they do not have the data. I think there are a lot of advantages as well for SBA in terms of being forward-looking and explicitly considering ENID.

Moderator: One final online question: who is the best person to do all the validation work?

Mr Kelliher: It is interesting because most of the validation I perform has been part of a second-line validation team, but in the Basel framework the third line does the validation because the second line is more involved in the risk modelling. The best person is the one with the best understanding of the risk profile. It is always good to have an external person coming in every now and again to challenge, to get that freshness of approach. But within the organisation, the best person is the one who has the best understanding of the risk and of modelling. I am agnostic between second and third line, so long as they have the skills and understanding of the profile.

Moderator: We have one final question in the room.

Questioner: I just wondered how difficult it is for benchmarking in either model because we do not have sufficient data or sufficient data to benchmark it against. Even when we scenario-test them, what are we going to benchmark against?

Mr Kelliher: That is a very good question. There are some good surveys I find in terms of benchmarking methodology. ORIC do an annual capital management survey. KPMG's technical practices survey has got useful information on operational risk and they have also done talks at the Life Conference. EY and PwC have different surveys. Benchmarking gives you comfort that your methodologies may be comparable to peers and also sometimes highlights new developments. Recently, I have noticed that one or two life offices are starting to use BN, whereas before this none of them were. That gives you some indication of where modelling is going. Benchmarking also acts as a sense check in terms of the diversification benefits and the actual level of operational risk relative to peers, but I would still take it with a pinch of salt. You could be comparable to peers in terms of your operational risk as a percentage of total risk capital, but having similar operational risk capital requirements doesn't mean much if you have a greater operational risk exposure and a weaker model. The fact that you are in line with your peers does not then mean much.

Moderator: I now bring the meeting to a close. It has been a fascinating presentation. Thank you also for all the great questions that have been posed.