
Good practice guide to setting inputs for operational risk models

Abstract of the Edinburgh Discussion

[Institute and Faculty of Actuaries, Sessional Research Event, Edinburgh, 6 March 2017]

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. and Cantle, N.J. Good practice guide to setting inputs for operational risk models. *British Actuarial Journal*. doi: <https://doi.org/10.1017/S1357321716000179>

The Chairman (Mr A. H. Watson, F.F.A.): Tonight’s speaker is Patrick Kelliher. Patrick is managing director of Crystal Risk Consulting. He has worked in life insurance for over 25 years and has extensive experience of operational risk analysis and modelling. As well as being a Fellow of the profession, he also holds the Chartered Enterprise Risk Actuary qualification in risk management.

Aside from chairing the Operational Risk Working Party, he is a member of the profession’s Risk Thought Leadership sub-committee and has produced numerous articles and papers on risk topics down through the years.

Mr P. O. J. Kelliher, F.I.A.: Before I start on the topic of inputs to operational risk models, I would like to set the scene by making a few observations about operational risk in general.

The first observation about operational risk is that it is a very diverse risk category. It covers everything from cyberattacks to mis-selling; from discrimination cases to processing errors. Under Basel II, there are seven high Level 1 categories of operational risk and 20 Level 2 categories of operational risk. I led a working party that looked into risk classification and we identified over 350 sub-types of operational risk.

The second observation is that there are a lot of similarities between general insurance and operational risk. For one thing, there is the potential for catastrophic losses like the £17 billion that Lloyds Banking Group (LBG) has incurred to date on Payment Protection Insurance (PPI) mis-selling. Like general insurance, there is a problem with latent exposure, what general insurers call incurred but not reported claims. An example of this from the operational risk space would be the losses suffered by insurers in the early “noughties” in respect of mortgage endowments sold 10 or 15 years earlier.

And even where we are aware of an issue that has come to light, the cost can be uncertain – what general insurers call their reported but not settled risk. To take an example again, looking back, the LBG PPI mis-selling provision started off as £3.2 billion in 2011. It has since grown quite considerably.

The difference between operational risk and general insurance is that general insurers actively go out to take on general insurance risk, whereas operational risk is, generally, not something for which you would have an appetite. It is something that happens as a consequence of doing business. We can mitigate it through the use of controls, and by investment in improving our risk control framework in insurance and the like. But, ultimately, there will always be some residual operational risk which we cannot remove.

Operational risk will evolve with our business model so, as we enter into new markets, our operational risk profile will evolve. An example of that is life insurance companies starting to offer a wrap proposition. That would generate a whole set of new operational risks which you need to take into account.

There are also technological changes in operational risk. Cybercrime is an increasingly important source of operational loss.

There can be changes in the legal environment. I cite an example of the European Court of Justice rulings on overtime pay which has cost some firms quite a lot of money.

I would say that, almost by definition, operational risk will evolve faster than the general insurance market can come up with solutions to mitigate it.

I would also like to set the scene in terms of the operational risk framework. Operational risk models are only as good as the framework with which they interact. Without a robust operational risk framework for identifying operational risks, our modelling results will not be worth anything.

A particular issue that I have come across in operational risk modelling is taxonomy, the categorisation of risks. I think that there is a need for a very detailed operational risk language to define what exactly is covered under what risk. If you do not have that very granular taxonomy, if it is vague, then you do risk operational losses being mis-categorised. I have been involved in many scenario analysis exercises where the same kind of risk has come up under two different categories because people were confused as to where the risk was covered.

The thing I find very scary is when you get the situation where we are missing key risks because people assume that they are covered somewhere else. Detailed taxonomy is crucial to understanding your operational risks.

Now, launching into the inputs themselves, obviously a key source of input to operational risk models will be internal loss data – a natural starting point for operational risk modelling. Indeed the Basel operational risk framework places a lot of emphasis on internal loss data.

There are a lot of problems with internal loss data. Most firms only started collecting operational risk loss data in 2000 or so. There is a dearth of data before then. Again, using the general insurance analogy, it is like trying to model windstorm losses when you do not have sight of any kind of experience before the year 2000, you would be missing out on the kind of events such as Hurricane Charlie in 1987. So there is a problem with high-impact low-frequency events which are not in the data – sometimes termed ENID (Events Not In Data) under Solvency II.

The other problem with internal loss data is the fact that it is essentially a retrospective view. There is a big question as to the relevance of some of the data. For many life offices, a large part of their loss data will refer to pensions and mortgage endowment selling claims. Looking forward, a lot of

pensions business will have already been reviewed and mortgage endowments will be time barred, so the actual prospect of future mis-selling claims from those particular sources is quite remote. That particular chunk of data is not that relevant any more.

At the same time, you might have entered new markets such as wrap. Entering the bulk buyout market will bring in new sources of risk which are yet to emerge in the operational risk loss data.

In terms of challenges that the working party identified, which we would like to highlight, one challenge of internal loss data is loss impacts and whether they are relevant for capital modelling.

You will have things such as fines and compensation. They will have a direct impact on own funds, but a lot of the time internal loss data will also pick up things such as loss of future sales or reputational impacts. It is important that we collect those things for operational risk management practices. However, our own funds do not allow, generally, for the value of future new business sales so it would be inappropriate to include that in your loss data for modelling of risks to own funds.

Another issue we thought worthwhile noting is the exposure to third parties such as fund managers and outsourcers. A lot of errors affecting your business can be made by third parties. But they may not show up in internal loss data because the third party often indemnifies the loss. Our view is that you should try as much as possible to capture those losses, even if they are indemnified, because, in extremis, the third party may not be in a position to indemnify those losses.

I think that there is a need to understand exposure to third party errors before we allow for indemnification.

As we went through the literature we found a lot of really good material, particularly from the Basel Committee on Banking Standards, which should be considered not just by banks but also by insurers. There is a particularly good paper, BCBS 196, on AMA approach to operational risks modelling which I found was very good on internal loss data.

Moving to external loss data, many funds are seeking to bulk out their loss data with losses of their peers. They might use confidential loss data sharing services like Operational Riskdata eXchange Association or Operational Risk Consortium, or they might use something like Algo First, which seeks to collate publicly available information.

There are a number of issues with external loss data. Just because you extend the number of firms in the data set does not necessarily mean that you are going to address the problem with low-frequency high-impact events. What you tend to see is everybody having the same latent exposures and then everybody being hit around the same time. For instance, take the example of PPI mis-selling. That was a latent exposure in the mid to late noughties for banks, but that only started showing in loss experience from about 2010–2011 onwards. If you were modelling operational risk in the noughties, even using external data would not have told you anything about that particular exposure.

Another problem with external loss data is that a lot of the time it is categorised differently from the categories you use for modelling operational risks. There are questions regarding mapping external loss data to your own categories.

There is also the need to scale the losses to your own particular business: I mentioned the £17 billion that LBG lost. That would be inconceivable for a small building society with £1–£2 billion in assets.

Conversely, you could have a small shock in a small friendly society and that could be a large part of that society's assets and, if replicated, could be quite severe.

Another issue with external loss data is relevance. I mentioned we have the problem of relevance with historic loss data internally in terms of changes in the risk profile, and in terms of changes in risk controls. When you bring in external data you are also bringing in data from firms with different control framework strengths and different business models. The losses that an annuity writer might incur might not be the same, nor as relevant, for an asset manager, for example.

There are quite a lot of problems with integrating external loss data into operational risk models, but even if you cannot fully integrate external loss data in the models, it is still quite useful to modelling in a number of ways. The first one is as an aid for scenario analysis, which we will come onto later.

The next would be to use it in terms of validation of model results. Finally, external loss data are essential for "business as usual" risk management: it is always much better to learn from other people's mistakes rather than to experience them yourself.

Turning now to scenario analysis, in terms of the literature review of the working party, one thing that kept coming up was that historic loss data on its own is not really suitable for operational risk models. It needs to be supplemented by a prospective view of risks. Generally, that prospective view is delivered by scenario analysis.

To my mind, scenario analysis is a key input to operational risk modelling. If it is done right, you can capture very low frequency, high-impact exposures that may be missing from the loss data. You can also reflect changes in your business model, new product launches and changes in controls, etc.

A big problem is, intrinsically, scenario analysis is subjective. There are many problems of bias. One bias may be where scenarios are anchored on past experience. Rather than looking forward in terms of what could go wrong in the future, the scenario tends to be anchored with what has gone wrong in the past. In that case, you really lose that forward-looking view that scenario analysis can give you.

Another problem is where scenario analysis results might be rejected if we find them uncomfortable, for instance, if you had a scenario analysis result which was materially greater than a standard formula allowance. My view on that is, even if it is quite large, we shouldn't discard it. If you think it plausible, you need to factor it into your modelling. The way I look at it is that it is one thing to fool the regulator but, once you start fooling yourself about your own underlying exposures, then you are in real trouble.

It is important to stress to people that there are no holds barred in scenario analysis and to encourage an open and frank discussion of what the potential exposures might be, and then take it from there in terms of modelling.

In terms of the requirements for scenario analysis, we have identified a number of key requirements for scenario analysis. The first is detailed preparation. There is a need to identify subject matter experts throughout the business who can contribute to scenario analysis. They need to be supplied with very detailed background material: things such as historic loss data, both internal and external; the results of risk and control self-assessments; and details of any kind of new product lines and business model developments which might have impact on exposure.

Having given that information and run the workshop it is important to set time aside in the process to allow for follow-up analysis. All too often people try to derive a loss figure in the workshop. Usually, it is very crude. You need to set aside time to do proper research of any loss impacts identified.

Being subjective, a crucial aspect of scenario analysis is independent review and challenge. Under Basel this is a role earmarked for the internal audit third line but I have participated in second line reviews of scenario analysis exercises. The key thing is that whoever is doing the review and challenge should be independent of the people who are running the workshops.

Documentation is another crucial stage in scenario analysis. This should cover not just the particular scenarios we selected for our capital calculations, but also the scenarios that we rejected and the reasons why we rejected them. That can really help in terms of demonstrating the breadth of your scenario analysis and coverage of all the risks.

Finally, there needs to be executive ownership and sign-off. Ideally, I think you would have an individual executive owner responsible for the assessment of a particular category. For instance, employee relations losses might be the remit of the HR director.

They should take an interest in making sure that the final figure is reasonable. In that way, by getting executives to own a particular scenario analysis, the insurer obtains appropriate buy-in throughout the business.

Correlation assumptions are obviously a key determinant of operational risk capital. The view of the working party is that empirical correlations will probably not be satisfactory. I have seen examples of correlations where you have some very strange results caused by a lack of data, and also the fact that you might be missing low-frequency high-impact events means you would not capture any tail events. So, it is necessary to supplement any empirical calculation based on historic losses with expert judgement.

The trick with that is how you elicit that expert judgement. If you are modelling operational risk by Basel II Level 2 Risk Types, that comprises 20 different operational risk capital figures. You will need to set 190 assumptions for the matrix. It is impractical to set 190 assumptions, let alone review, challenge and secure executive sign-off for such an amount. There is obviously a challenge here.

One possible solution may be to consider correlations between blocks of businesses, for instance, to group operational risks by Basel Level 1 categories and consider correlations that way.

Another idea we had was that you could consider the impact of certain causal factors on individual risk loss assessments. You could consider whether a particular operational risk would change if we had high staff turnover, or what the impact of a flu pandemic on this particular loss category would be.

By considering the impacts of these causal factors on individual risk categories, you might be able to elicit something about common dependencies between operational risk categories.

Whichever way we approach this activity, given the subjectivity that is likely to be involved, there is a need to have extensive independent review and challenge, and ideally executive sign-off, given that it is of such key importance to an operational risk model.

In terms of other inputs, there are other key aspects to consider, such as risk mitigation, that is, allowances for insurance or allowances for indemnification of losses by outsourcers.

The first point that I would make on this topic is that it is quite important to understand where you might have implicit, as well as explicit, allowances for insurance in your operational risk model.

Just to explain that point, what I have often seen with business continuity scenarios is that they tend to discard scenarios that are covered by insurance. You are not obtaining a good picture of your gross exposure as people are implicitly allowing for insurance cover.

If you are allowing for insurance, I think you need to go through that insurance policy to understand in detail what is covered and what is not covered. An example would be if you are modelling cybercrime.

Cybercrime policies will generally not cover fines. When you are modelling cybercrime operational losses, you have to know what the fine component is since it will not be covered by an insurance policy.

Similarly, with outsourcing, when making assumptions about the ability to reclaim losses from the outsourcer it is important that we do the research to ensure that we can claim back those losses under the outsourcing agreement and that they have the financial strength to honour indemnities, particularly when dealing with extreme losses.

Another issue is allocation of operational risk capital. Often in modelling operational risk, you will model at a very high level, for example, at country level, and then there is a need to allocate back capital to individual legal entities. There is a need for research to support whatever approach you use for allocation. For example, I have come across instances where service level agreements precluded employment relations losses being charged back to subsidiaries. That would obviously have a bearing on the element of employment relations losses allocated to a particular legal entity. You need to do a lot of research into whether you can charge a loss back to a legal entity.

In UK life insurance, with-profits is a key consideration. There might be scope to charge operational risk losses to the with-profits fund and to with-profits policyholders. Whatever the assumption, it needs to meet the requirements of the Principles and Practices of Financial Management (PPFM) and it needs to be consistent with Treating Customers Fairly. Good practice here would be for any assumption regarding allocations of losses to with-profits funds to be reviewed and approved by the with-profits actuary and the with-profits committee.

To conclude on the operational risk working party key findings: in terms of our literature review, loss data on its own is not going to be good enough. It needs to be supplemented by a prospective view, particularly scenario analysis. There are some good examples of good practice out there. I have mentioned the BCBS 196 paper on loss data.

A prerequisite of any kind of operational risk model is a good operational risk framework. In particular, there needs to be detailed taxonomy to avoid mis-categorisation of losses.

On internal losses, there are issues with low frequency, high-impact ENID, as well as the relevance of loss data. One thing to which we would draw attention is understanding what loss impacts are included in the loss data and ensuring that they are actually relevant in the modelling of funds.

Also, we would highlight that on a gross basis you might have significant losses being incurred by fund managers and outsourcers, but that might not be coming through in your loss data because of indemnification. There is a need to get a handle on just how big is the exposure.

On external loss data, even extending your data set to include other firms, might not fully address the issue of low-frequency high-impact events. There are obviously going to be issues with the relevance of different businesses, different business models, different products and different control framework strengths. There is going to be an issue with scaling. It is going to be very challenging to incorporate external loss data into operational risk models. If we cannot bring them directly into models, however, we can consider these for scenario analysis, model validation and risk management.

Scenario analysis is a key component of operational risk modelling. It needs to be done with appropriate rigour. There needs to be extensive preparation. The process needs to allow for plenty of time to follow-up scenario assessments and further research. There needs to be independent review and challenge. Documentation of results, including scenarios discarded, is needed. Then there needs to be executive ownership and sign-off.

With regard to correlations, to reiterate, we think that empirical correlations are likely to be flawed and that there is going to be a need for expert judgement. There will be issues in eliciting that judgement.

Other assumptions in terms of risk mitigation and in terms of allocation of capital to legal entities need to be supported by extensive research into things such as insurance policy cover and, for with-profits, whether a particular chargeback is allowed in the PPFM.

Mr A. J. Clarkson, F.F.A. (opening the discussion): I agree with a lot of what Patrick said. My experience is that there is a large amount of judgement involved in the assessment of operational risk capital. It is quite plausible to come up with markedly different results from the changing of one plausible assumption to another. Judgement around inputs into scenario testing is hard enough; judgements around correlation assumptions are even more challenging. In that context, it is really important that the key assumptions that drive the overall capital amount and the sensitivity of the result to those key assumptions are made visible to the Board.

The paper does not focus on the use of operational capital amounts within the business. I think it would have been an interesting area in order to obtain an insight into best practice.

For example, do companies take the capital impact into account when taking key decisions such as headcount reductions or new proposition launches?

How do companies relate operational risk capital to risk appetites and the day-to-day operational risk framework?

To what extent do companies use key risk indicators to adjust capital amounts during the year, and what are effective key risk indicators?

This might be a potential area of future work for the working party.

I have a few specific comments on the sections in the paper that consider data. Section 5.1.3 talks about change risk and the need to allow for cost overruns in change projects. An alternative place to

allow for that is within expense risk. The key is making sure that it is somewhere, not specifically which category you put it into.

Section 5.1.6 talks about obtaining data from third parties. While I agree in theory with that, I wonder how realistic that is. The paper suggests that new agreements should require such reporting. In reality, something may have to be given up in negotiations to obtain this and a balance will have to be struck.

Section 6.2 notes that one of the three key challenges with external data is the need to reflect any changes in control in the external company since the losses occurred. I would have thought that a much bigger challenge would be that the control environment for the external company might differ significantly from that in the company itself which could render the data relatively meaningless. Patrick (Kelliher) mentioned this point.

Section 8 then discusses scenario analysis, and the paper highlights that a common failing is where the same risk is considered under two separate categories or where a material risk is missed. Again, Patrick mentioned that in his opening remarks. In my experience, this is indeed a key challenge. An important mitigant to that is ensuring that someone is tasked with reviewing all the proposed scenarios to ensure that is not the case.

I agree with Patrick, and the authors of the paper, that there is a role for risk in the independent review of the proposed scenarios. The key skills and the knowledge required for such a review are a good understanding of the business environment and context, the control environment, and the emerging external environment. These are key skills that should exist within Risk. It is perfectly possible to involve people within Risk who have had no involvement in the facilitation of any scenario workshops, thus ensuring appropriate independence. Again, Patrick mentioned this as a key point.

Another key challenge is how to ensure that a scenario used for a particular risk category produces a result that is representative of all possible scenarios within that risk category. I think it would have been interesting if the paper had discussed that challenge. It is very challenging for a subject matter expert to judge the frequency and potential impact of low-frequency high-risk events as they, by definition, have little past experience with which to frame their judgements. It is even more difficult for them to judge the extent to which input assumptions should be adjusted to reflect the fact that there are a number of potential scenarios that could be applicable within a particular risk category.

Finally, a key aspect of the validation of operational risk capital is a top down, reasonableness assessment. How does the amount of capital compare with the previous assessment, and does that make sense in the context of any changes in the control environment, headcount changes, the emerging regulatory environment, and the level of operational change and stretch?

Mr Kelliher: I think the use test is one area into which we need to look. It is important for models. The operational risk working party is looking to produce a series of papers. I will take a note of the use test.

One thing I would say is, rather than the actual capital figure being important, sometimes it is the process that you arrive at this which is the real benefit for the firm. Just trying to capture losses and also doing scenario analysis, can help a firm's risk management, even if you are not using an internal risk model. The insights that you can gain from doing scenario analysis can really help.

I agree with your point regarding change risk potentially being expense risk. As you say, the key thing is not so much whether it is expense risk or operational risk it is just the fact that it is covered somewhere.

As far as third parties are concerned, I think we were very cognisant that it may be very difficult to obtain details from third parties on control failures and the like. However, a lot of the time there will be certain reports that are already being produced and which should be circulated more widely, for example, on fund management breaches.

With outsourcer loss, a lot of times that will come through your complaints department. There will be compensation paid by your firm. You then recover that from the outsourcer.

A few points on representative scenarios. I think it is a key challenge. I do not think that it is plausible to go through 350 or so different operational risks and come up with a scenario for each of those for executive review, challenge and sign-off. We are necessarily going to have to look at representative scenarios.

In terms of approaches I have seen in practice, one approach would be to first try to consider the frequency of a material loss arising within a category, perhaps by reference to fixed frequency bands such as 1-in-2 years, 1-in-5 years and so on. Next, try to come up with some typical and severe loss examples, so for business continuity risk the typical loss could be based on, say, water escape while an extreme loss would be based on a terrorist bomb. This could then be the basis for modelling.

I agree that top-down validation is needed to check on model results to ensure that they are reasonable.

Mr J. R. Crispin, F.F.A: We went through this exercise last year. We have taken it to the Board. I agree that scenario analysis is a key part of the process. One of the considerations we were taking is how you get the intercept between what you are doing in scenario analysis and what you do from all the underlying data. We have lots of data on small risks at high frequency. We have not so much data on low-frequency, high-impact risks.

The approach we took was to focus much more using scenario analysis at the high-impact end of the spectrum, and within that to focus on making sure we were using generic scenarios. So, we had fewer what we call “better” scenarios to try to capture the point you were just making around how you address the different frequencies.

I note that you have not said anything about how you would then fit the distributions, or about what distributions you would fit to the scenarios, and also how you combine frequency of scenarios with the distribution scenarios.

I think the point you are making is that a structured process is key, and that the documentation is clearly very important in obtaining approval from both the Board and then the regulator, if you are willing to have an internal model.

Mr Kelliher: In terms of the types of distributions, I purposely did not try to delve into this area.

From my point of view, what I have seen a lot from operational risk modelling is that too often actuaries jump in and obtain a lot of data on very small scale losses, fit a distribution to these and, lo and behold, you have a certain capital figure.

I have heard of a case where an insurer developed a very sophisticated distribution based on the small loss data that they had and they came up with a capital figure of £50 million. But then came the question: “But did you not pay £100 million 2 years ago for mis-selling losses”? In this case, what happened was they purposely excluded a large mis-selling loss from the capital figure, which then looked a bit light and indeed it was because they did suffer a £200 million loss a few years later.

So we purposely did not try to delve into this area. The focus of the paper was not so much on the stats and maths; it was more putting the framework in place. The key thing with operational risk is putting the basics in place in terms of collecting enough data, but also carrying out scenario analysis. With a structured process to think about the operational risk exposures, you can then review and challenge effectively.

Mr J. E. Gill, F.F.A.: I want to make a couple of points. If you are genuinely frightened by a credible scenario, then why are you in that business line in the first place?

A good example of that is if you look at things like client money and client assets and apply the scale of fines available to the Financial Conduct Authority (FCA). If you multiply very small numbers by very big numbers you end up with a quite significant potential loss which encourages businesses to have the right kind of mitigants in place, given the appropriate scenarios.

Second, I am not sure where the working party have got to on this: the historic cyclicity of operational risk losses. There is probably more evidence of this in banking than there is on the insurance side. You could argue that the typical operational risk actual losses are benign and pretty straightforward for 9 years and then in the 10th year there is a blow-out. We have seen that pattern repeat itself several times.

By definition that usually means that whatever you have set as your risk capital for 9 years has been far in excess of what is required, and when you really need it, it is not big enough.

I wonder whether the working party have thought about how to address the overall cyclicity issue that is at the heart of risk capital quantification.

Mr Kelliher: On the cyclicity point, during the boom times we had a huge surge in mortgage lending by the banks, but it was fundamentally unprofitable. Because of that there was a scramble to sell ancillary products like PPI.

Then the next thing was that crazy mortgage lending ran out of steam and there was the crash. Then the PPI mis-selling came to light. It is not just PPI mis-selling. In the boom times, there was a lot of London Inter-bank Offered Rate (LIBOR) fixing, there was a lot of mis-selling of sub-prime securities.

So, the chickens come home to roost. It is not just from the current financial crisis, if you look back to 1929 there was a lot of major fraud emerging when the bubble burst in 1929. History does repeat itself – but not exactly.

In terms of the paper and cyclicity, it is not in the paper but the working party did separately consider some recent changes in the Basel framework on moving to a new standardised approach. I think we did make some comment about the length of data used in calculations. How do you address this problem of cyclical events?

A particular issue they had with the old system of Basel was that it was based on gross income. What happened was the gross income started contracting shortly after the downturn which obviously was when all the losses started to come through. There was the question: how do you address that aspect?

The Chairman: Something was said about bulk buying of annuities. How could we extrapolate this into the world of pensions? Should trustees consider doing something around this area?

Mr Kelliher: That is interesting. One of the topics we are considering is operational risk in pension schemes.

In terms of the evolution of operational risk, in the 90s and early “noughties” huge losses were arising in life insurance with pensions mis-selling and with mortgage endowments.

Since then there has been a big wave of losses in banks: PPI mis-selling; London Inter-bank Offered Rate; and so on. Maybe the next wave of loss is going to emerge in pension schemes. I think The Pensions Regulator has already called out issues with defective record-keeping and problems with valuations. For instance, you have pension scheme debts on the balance sheet. If the actuary gets that wrong, you could find the deficit widening.

There are risks for the scheme sponsor in terms of widening deficits as a result of operational risk. Also, there is going to be operational risk for actuaries advising on schemes. So, yes, watch this space.

The Chairman: Maybe that discussion on pensions has helped some of the pensions actuaries here to formulate some thoughts.

Mr A. C. Martin, F.F.A: There is the very recent report from the National Audit Office (NAO), who reported on the extra payments of some £770 million of commutation amounts paid to firefighters and police officers.

The NAO reported on the late payment to some 34,000 officers in uniform of their commutation that should have been paid on an actuarial equivalent basis from the turn of the century to 2006. Internal emails between the Treasury, Department for Work and Pensions and the Government Actuary’s Department are excellent reading if you are looking for a case study. We, as taxpayers, simply paid out late what should have been paid an awful lot earlier.

More for the future, the obvious operational risk is probably payments from defined benefit (DB) schemes to defined contribution (DC) arrangements, where even the former Chancellor recognised the potential for mis-selling and required financial advice to be provided for amounts over £30,000.

Advice is required but can, of course, be ignored. Operational risk will therefore be for people paying out as well as the people receiving. There is at least one trustee in the room agonising over whether to ask the question as to whether advice has been ignored. The reality is that it is only the ones where advice has been ignored and things go wrong, perhaps only one in five or 10, that will come back to provide the litigation for the next round of mis-selling.

My question on operational risk generally, however, is how to put in the softer elements. I describe one of these as communication. The NAO remarked upon government departments not talking to

each other. I think the culpable challenge would be internal departments (e.g. economists, investment analysts, sales people) actually talking to each other to know it is now beyond rumour that some people just sign their name and have the rest of the form filled in by the sales person, who is remunerated for the extra business irrespective of whether it is in the customer's interests or not? Communication, I think, should be the real challenge.

Mr Kelliher: As you said, £770 million due to errors in late commutation factors which you need to compensate members for.

We may consider what would happen for a private sector scheme if a similar lawsuit was brought against the trustees which successfully claimed for underpayment. Suddenly, the next thing is you have a much bigger deficit in your pension scheme which is going to sit on your balance sheet. That could be a very interesting area where pension scheme operational risk affects the sponsors.

The area of DB to DC transfers is going to be a very difficult area. It is not going to impact the trustees but the independent financial advisers (IFAs) who are involved in giving advice. The life insurers who are going to accept a lot of these DC transfers could also potentially be affected.

On this insurers need to consider their provider responsibilities. There are a lot of responsibilities on a life insurer as a provider to provide the information necessary for people to make an informed decision. What I have seen of these DB to DC transfer exercises is that a lot of the time you have a critical yield which covers the investment risk, but you do not have much on the longevity risk. There is potential exposure for life insurers in accepting these transfers if they are deemed not to give enough information on this risk.

Even where you do give advice you can have insistent customers. There was something from the FCA quite recently about insistent customers. A lot of the time, even when the advice was not to proceed, they still proceeded. You think that surely the IFA must be in the clear, having told them not to proceed, but sometimes what you find is that, when the FCA looks back at the recommendation not to proceed, if they feel there was any gap in the analysis, then they could maybe turn around and say that you did not do enough to make clear this particular option. For that reason, even the fact that you have told them not to do it may not be enough to protect you.

Mr R.R. Ainslie, F.F.A.: When I worked with a term assurance and critical illness insurance provider one of the main operational risks that we faced was administration error. I do not mean the kind of administration where you collect a direct debit twice or something like that. I am talking about the kind of risk where you perhaps give someone who is terminally ill a £10 million term assurance. You make a mistake in the underwriting process or the new business keying process.

One of the very useful sources of data we found for that was the risk control process itself, which was checking not all cases but some cases. You would do sampling of the work that is being done to, it is hoped, identify errors and put them right before they become cases.

But it is a very useful source of data that tells you what your future loss exposure might be from cases going wrong, should the customer come to claim.

The interesting challenge we had with that was incorporating it into the pricing of the business, which of course we should do for operational risk. Operational risk is a risk that is within our

appetite. You should be looking to quantify it, price it and possibly include it as a decrement within your pricing models, which is what we looked to do, and have that sitting along there as a cost that was charged to customers. This is slightly unfortunate for them as the average customer pays for the operational errors, but that is the way it is.

We found that process of collecting data and then feeding the assumptions into the pricing model was a very useful way of making sure that operational risk was properly charged for.

The trickiest problem we had was when the losses began to come through. That was costing us as a business. We had to go through the process of explaining to our colleagues in finance, and I must confess that we probably should have set the ground more soundly beforehand. Explaining that even if you paid £2 million or £3 million in claims that you probably should not have because of an error, that is not a problem because the pricing in your modelling assumed that you were going to be paying £4 million.

Entering into these problem discussions and explaining that a loss is not really a loss but your operational losses could be contributing experience profits was interesting and, perhaps, an example of the use test that Alistair picked up.

The other advantage was by incorporating operational risk into the pricing models you could encourage wider parts of the business to engage with how to better manage operational risks. They would filter through into improving the prices and improving the competitiveness of the business.

Mr Kelliher: In the area of protection I think that operational risk errors can be quite crucial. What you often find when you have that underwriting error is not only have you accepted somebody in poor health on normal terms, for example, but you will find also that the reinsurer will not cover it. That is an additional risk element.

It is probably inevitable that there will be some degree of underwriting process error. I think somebody pricing the business without allowance for that is probably under-pricing that particular protection product.

Mr H. R. D. Taylor, F.F.A.: On the topic of pensions, just thinking about emerging risks rather than risks which are extant at the moment, it seems to me that wherever you have major change in the way a market works, there is likelihood that new risks will emerge. One of the biggest changes, particularly in the DC pensions market in recent years, has been pension freedoms. That, over the next 5–10 years, is going to lead to a massive increase in the number of individuals who are exercising their right to do pensions drawdown or to take lump sums.

Whenever you have large amounts of money in circulation, as banks know, there is potential for fraud. Fraudsters tend to be stopped and then find other ways of accessing pools of money that are moving.

Given the existence of large pools of money that may be moving over the next 5–10 years for individuals who may not be financially sophisticated, fraud is a huge area of risk. I am not entirely convinced that the pensions industry is doing quite enough to protect these individuals. With DC it is the individual who bears the brunt of an event that goes wrong.

The other thing is around an individual's understanding of what they are likely to obtain from a pensions pot if they do drawdown rather than buying an annuity. We know that there are still some annuity sales but what is interesting about the mass-market that is going to emerge with pension freedoms is that the individuals are relatively unsophisticated. They either cannot afford, or do not want, to spend money on detailed advice. Therefore, they are relying on a mixture of guidance and also the emerging artificial intelligence or robot advice processes that will be in place.

I think that all of these things are potential sources of risk for everyone who is involved in DC pensions.

Finally, just to pick up on an earlier point of transfers from DB to DC, a point of leakage of individuals' apparently secure benefits in DB is the attraction of getting control of the DB assets. They are then transferred into DC. There is a massive transfer of risk around their long-term income from the DB scheme to a DC scheme, and of course any of the other risks that exist in money moving around the DC market immediately become a feature of life for them. So, I think that there is plenty to be looking at.

Mr Kelliher: It is not just pension schemes that will be affected by some of these issues. There will also impacts on the life insurance providers who will be servicing these drawdowns. I suspect with a lot of them there will not be giving any advice. But then there will be a lot of attention if something goes wrong around how they were sold. Was the literature that was provided clear enough about the risks? We must be very careful.

Regarding fraud risk, I think that we are already starting to see an uptick in fraud. There was always a small level of fraud but pensions unlocking means we are now starting to see fraudsters moving into this market.

I was at a very interesting presentation last week on cyber risk and fraud. It is quite scary the sophistication of fraudsters and the operations that they have. They have their own hierarchy; they are businesses in their own right. They will be out there trying to gain information on people with pension pots, looking to relieve them of their money.

Mr J. M. Black, F.F.A.: When you receive the report on operational risk, there is always this feeling that you can control it, and if you do all the things that you are supposed to do, you should see it going down from 1 year to the next.

With that context in mind, and given the discussion which seems to focus a lot on mis-selling risk, do you think that we, as an industry, are becoming any better at understanding and mitigating mis-selling risk?

Mr Kelliher: I think yes, we have. Just the basic things. You would not find many life offices now that would appoint banks' sales staff as appointed representatives which they used to back in the early 90s. I think that we are getting smarter about the mis-selling risk.

The problems with mis-selling risk are changing. While we might no longer be dealing with direct sales, I think what could be a really significant thing is provider responsibility. We could see that the advisers are no longer our appointed representatives and we may think we have washed our hands of

mis-selling. That might not work going forward. What might happen is that there could be some comeback from advisers. They might say you did not explain that particular product to me so this is the reason I was mis-sold.

A general point on control of risk, the one thing that struck me from the meeting last week on cyber risk is the sophistication of some people out there, the cyber criminals. It is quite awe-inspiring. Their ability to hack your systems is quite phenomenal.

One classic case in financial services is a health insurer in the US called Anthem. Hackers stole not just current records but also legacy records. All told, they stole nearly 80 million people's records. I do not think that attack was due to fraudsters, I think it was state-sponsored, but I think that what we can say is you can obviously invest in controls up to a certain point but you are still going to be facing the danger that at some stage a fraudster will get through.

Looking at the changes in data protection risk, beforehand the Information Commissioner's Office could fine you up to about half a million pounds for a breach. That is changing. With the General Data Protection Regulation they can fine up to 4% of turnover.

We are becoming smarter but I think operational risk is moving on. Some risks morph into different kinds of things like mis-selling. New risks come along like cyber. I think that we are always going to be up against it.

Mr D. G. Robinson, F.F.A.: I wonder if you would like to say a little more about scenario planning. We have teams of executives sitting round a table. They get can become very boxed in, in terms of their thinking. How do you help them out of that box?

That is really a follow-on to the point made by Jim (Black) around mis-selling and the point about PPI mis-selling in particular, where the operational risk now is that you do not reserve enough for mis-selling claims.

The original operational risk was that they were going to be found out. So, in your taxonomy of 356 classes of operational risk, was there a risk of being found out? I think it is less of a risk nowadays, as you say – the behaviour of product providers, and so on. But, sometimes, businesses do not realise that they will be perceived to be doing wrong in the future.

So how do you get out of that difficulty?

Mr Kelliher: I think that is a general problem. Groupthink is a particular problem. I would not say that the issue was whether we are going to be found out. I think there was probably a blindness among banks to potential losses that they could incur from PPI mis-selling.

I am still amazed that when you look back to 2005 and you see the Office of Fair Trading reports saying PPI was a bad product, it had very high commission, and it was not a good deal for the consumer yet banks were still selling it. I am always stunned why they didn't think they might have a mis-selling issue though. I suppose 20–20 hindsight is a great thing.

There is a need to think outside the box. One matter which might have affected their thinking was when you started putting figures to PPI mis-selling; LBG at one stage were looking at one-sixth of its

profits. If they had to refund one-sixth of our profits in a year, that was probably on a par, if not greater, than their standard formula operational risk charge.

Maybe you have this kind of cognitive dissonance when you derive a figure that is so big you just put it to one side and say surely that can never happen because we have £2 billion of operational risk capital. Surely, we cannot have a loss more than that.

Mr Robinson: They were perfectly well aware at the time that they were doing wrong. When I was involved with the Association of British Insurers, probably 15–20 years ago, I tried to get them to take action against product providers who were paying 95% commission plus profit share, so only 5 p in every pound that the customers were paying was going to pay for the risk.

I do not want to go into personalities, but they had the answer that you are giving in a sense, which was that we are making too much money so we cannot stop and everyone else is doing it.

So the risk was there for many years. The chickens – or whatever – were going to come home to roost. That was going to happen. I think that they were blind to it. There was groupthink going on there, but they could have avoided it all 15–20 years earlier, if they had listened.

Mr Kelliher: This links with a wider aspect. A lot of people knew what was going on. If you take Halifax Bank of Scotland and commercial lending, a lot of people knew that the expansion of that particular book was crazy; but there were issues in terms of challenging the expansion. That is a general problem, not just with operational risk. There can be cultural issues when something is doing really well which prevent people from asking if it is not sustainable to make so much profit from customers.

You would like to think people would stand up; and I like to think that actuaries – we have a code of ethics – would be the people who would cry a halt. But that did not happen with the banks.

Mr Taylor (closing the discussion): Just picking up on a couple of points in that discussion, one of the big risks in any business is that the goals change. The nature of the game changes and it happens quite suddenly.

What happened with the banks and PPI was, fundamentally, the attitude to ripping off consumers changed, and the appetite for regulatory intervention on behalf of consumers is quite different now across financial services from what it was 20 years ago. I am sure what has gone on now was never envisaged 20 years ago. That is one big change.

Again, you would only detect that if, as David Robinson said, you have a group of executives who are prepared to think outside the box.

It is like running any business where there is disruptive innovation, your entire business model can collapse in a matter of years unless you plan to do something about it. The world is littered with companies like Kodak.

I guess that the final thing is that if there is something that is exposed, then there is a whole industry now prepared to jump in and take a share of whatever compensation they can obtain from clients, from insurers, from banks or anyone else.

That is probably one of the reasons why car insurance premiums are becoming so high, because of potentially dubious or even fraudulent whiplash claims. There is an entire industry now around generating these claims. Again, when you are in an ecosystem that is going to arbitrage against you, then there is potential for very large amounts of money to be shifted from the provider of the product to the consumers who are using the product.

Mr Kelliher: It is interesting that you mentioned ambulance chasers. This started with mortgage endowments. Although providers did not have to do a proactive review for mortgage endowments as they had to do for pensions mis-selling, which probably saved a lot of money, it led to ambulance chasers becoming involved with financial services mis-selling. They certainly realised that there was money in this mis-selling game. As soon as mortgages started to wind down, they could see where PPI was going, they got in on the act.

They are looking for other things. They are looking at State Earnings-Related Pension Scheme mis-selling. Things could change when we move to a flat rate pension, since how much people might be losing from contracting out will become more obvious. There might be a sea change in this area and maybe that could be the mis-selling scandal of the next 20 years.

The Chairman: I wondered if what was happening last week with the Lord Chancellor and the discount rate for Ogden Tables is an example of operational risk that some insurance companies have taken account of.

Mr Kelliher: I think this is something they should allow for. Maybe not as an operational risk, but certainly under general insurance risk we need to assess what could go wrong and how much will that cost us? I would be very surprised if general insurers did not have some scenarios for how badly it could hit them.

Maybe they were expecting not the full reduction, but I would be very surprised if general insurers did not have allowances either in operational risk or in general risk.

The Chairman: We have probably offered Patrick some ways to allow the working party to develop some of its thinking, including for us poor pensions actuaries. It sounds like there is going to be some work for the pensions research committee.