
A common risk classification system for the Actuarial Profession

Abstract of the Edinburgh discussion

[Institute and Faculty of Actuaries, 7 November 2011]

Contact

Patrick Kelliher, F.I.A., E-mail: patrick_oj_kelliher@yahoo.co.uk

This abstract relates to the following paper: Kelliher, P.O.J., Wilmot, D., Vij, J. and Klumpes, P.J.M. A common risk classification system for the Actuarial Profession. *British Actuarial Journal*, doi: 10.1017/S1357321712000293

The Chairman (Mr D. B. Martin, F.F.A.): Good evening, ladies and gentlemen. Welcome to tonight's discussion on a common risk classification system for the Actuarial Profession.

I would like to introduce one of the authors, Mr P Kelliher, FIA. Mr Kelliher is an Edinburgh-based actuary. He has worked for various employers but now has branched out on his own as a Solvency II expert.

Mr P. O. J. Kelliher, F.I.A. (introducing the paper): Coherent risk classification is the key to enterprise risk management (ERM). If you do not have a coherent risk classification throughout your organisation, you are going to end up with confused reporting of risks and allocation of responsibilities.

However, even if your own organisation's risk classification system is coherent, it is probably going to be different from those of other actuaries working in other organisations. This gives rise to the potential for confusion between actuaries. Therefore, the risk classification working party was set up with the idea of coming up with a common classification system that everybody could access, and the hope is that this will be a common reference point for discussion of risks between actuaries. We are not trying to come up with a definitive risk classification system. That would be too difficult. The aim is to give actuaries a common base with which to discuss risk.

I should like to touch on some examples of the multiplicity of risk classification systems and terminology in use and some of the potential problems caused by this terminology. I will also discuss some of the high-level principles which underpin this common risk classification and then I will go into the detail of the individual risks and the problems we faced when we were developing this system.

In respect of the multiplicity of systems in use, I thought I would give a few existing examples of high level categories. The first one is from the FSA's Systems and Controls Handbook. It has chapter headings for market, credit, insurance, liquidity and operational risk. It also has a group risk chapter, which relates to exposures to other parts of the same financial services group.

There is a similar classification system, which I understand BaFin, the German regulator, uses. It is in some ways similar in terms of market, credit, liquidity and operational risks; and underwriting risk, which I understand is close to insurance risk; but then it has a concentration risk element separate from credit risk. It also has explicit categories for strategy and reputational risk.

Turning to some practical examples in the private sector, Lloyd's Banking Group has disclosed high level categories of market, credit, insurance, and operational risk; but it also has the financial soundness category. This includes liquidity risk as well as tax, accounting and regulatory capital issues. Finally, it has a business risk category, which I understand broadly covers strategy-related risks.

Another example from the private sector is Prudential. Again, it has market, credit, insurance, liquidity, and operational risks. Then it has two types of strategy risk, one relating to its own business, and the other to the wider business environment and the risk it poses to its strategy.

I thought that the elements like market and credit risk were common to all until I came across this example from Aegon, which does not have market and credit risk. Instead it has investment and counterparty risk, which covers a mixture of market and credit risks. It also has mismatch risk, which covers other aspects of market risk, and operational risk and underwriting risk categories.

That was a just a brief illustration: even at a high level there is quite a large variance in terminology between different organisations. As you drill down to further levels you come across more and more areas of ambiguity and differences in terminology.

There are many areas of doubt, one of which would be non-disclosure. Some would class this as an element of fraud risk; others would class it as an aspect of their mortality experience.

Another is spread widening: is it a market risk or a credit risk? Is the shortage of buyers in the market a market risk or is it a liquidity risk? As we went through our work we came across more and more examples of potential confusion.

It is like the Tower of Babel in terms of everybody speaking their own particular language when it comes to risk.

Does it matter? I think that it does. There are two reasons why I would be concerned about actuaries using different terminology. The first reason is that the Actuarial Profession wants to move into wider enterprise risk management.

As part of this, we will need to get more and more involved in areas such as operational risk and strategic risk. I think these are the areas where there is the greatest scope for confusion between actuaries as this where there are the greatest differences between individual companies' terminology.

The second reason is that this potential confusion could be significant for when we try to collaborate on Solvency II.

I would give two particular examples for Solvency II. The first is mortality risk. I imagine most life insurers will have mortality risk as one of the risks on their risk register. When you drill down in terms of the internal model, you might have one aspect of this risk which you call parameter estimation risk in your firm but which might be called something else in another firm.

Also, changes in non-disclosure law can lead to a step change in mortality experience. Is that something that you would include in your mortality risk internal model or is it something you would cover under your operational risk internal model?

The second example relates to own risk and solvency assessment (ORSA). I think there will be a need to be able to break risk down to a granular level. For example, there are certain parts of equity risk that are mitigated by controls, like stock specific risks. Then there are other parts that we cannot diversify, like the beta risks. I think that offices, in terms of how they break these components down, will differ in terms of the terminology. I think that there is scope for confusion.

Moving on to the broad principles underpinning the common risk classification, our starting point was embedded value, comprising net assets and value in force (VIF).

Some obvious risks to embedded value are market risk, credit risk, insurance risk including the impact of persistency on VIF, and operational risk, which is a mixture of both immediate cash payments, such as compensation costs, and the loss of future income and VIF.

These address the amount of embedded value, or the amount of economic assets over liabilities; but, as important as the amount is the liquidity of assets relative to liabilities. So we need to have a category for liquidity risk.

Also, the embedded value is not the total worth of the company. There is an element of goodwill as well. This relates to the value of future new business streams and the various initiatives the management will take to increase efficiency. We therefore have a strategy risk category relating to the threats to goodwill.

Then we came up with a frictional risk category. The previous categories looked at economic impacts but often you will find, and Solvency II is a classic case, that changes in regulatory capital requirements can have a material impact on your business, even if there is no change in the economic risk profile. We have covered these risks, as well as taxation, in this frictional risk category.

Finally, tying it all together, we have an overarching aggregation category, which relates to the whole being greater than the sum of the risks.

The first key concept we decided to adopt in developing this risk classification was event-based classification. To give an example of what this means, our event-based classification would classify what happened to Northern Rock as a liquidity risk, a plain, old-fashioned run on a bank. Causal analysis might look into what caused that run. Was it poor governance? Was it flawed strategy?

Causal-based analysis is obviously important but one problem we found with a causal-based system is that often you can identify multiple causes. That is why we just focused on events.

The risk classification system focuses on gross risk and it generally excludes control failures except for operational risk. In particular, it treats asset/liability management (ALM) as a control, and focuses more on the underlying exposures of the liabilities and assets.

Reputation risk is covered under strategy risk. Reputation damage can affect withdrawals, but withdrawals come under persistency risk as part of the insurance risk category.

We generally include within market and other risks the regulatory and accounting impacts of these risks as well as the economic value impact.

I should like to touch on some of the issues that we came across in going through developing each particular risk category.

Starting with market risk, we worked from the FSA's definition, which is a very good one. It covers equities, properties, commodities, currency, etc. The one area in which we differ is that we have made a distinction between actual and implied inflation. Implied inflation is effectively the difference between the nominal and real yield curves. We have a separate risk category in terms of the actual RPI and NAE experience.

We have broken the market risk down into various components such as stock specific, sector specific, dividends, implied volatility and so on.

In coming up with market risk categories, we came across quite a number of demarcation issues and questions.

The first one relates to private equity. Is it a distinct category in itself or is it a component of equity risk? We opted for it being a component of equity risk on the basis that exit values are usually linked to stock market returns.

Another question was about interest rate risk. We have defined that in terms of changes in the risk-free curve. But what is risk free? Is it swaps? Is it gilts? We opted for swaps because this is where Solvency II appears to be going.

There is a need to distinguish between the liquidity effects relating to the balance of buyers and sellers in the market. This has an impact on the market risk but there is also a liquidity risk angle to this. We classed changes in the mid-market price of assets as market risk, while changes in bid/offer spreads relative to mid-market price would be a liquidity risk.

Another area of overlap is rogue trading losses. These are obviously market-linked. However, they are conditional on operational failings, so rogue trading losses are classed as operational risk.

Similarly, the impact of falling markets on new business sales comes under strategy risk.

We have 28 different categories of credit risk based on different types of exposure. These are broken down further by probability of default, exposure at default, loss given default, and migration risk, which relates to downgrades and changes in internal credit ratings.

In terms of the credit risk demarcation issues, probably the trickiest was how to distinguish between credit risk in terms of defaults and downgrades, and general spread movements under market risk.

A problem is that usually when a bond defaults or is downgraded, much of the default or downgrade will have already been factored into the bond price and hence the spread movement.

Another demarcation issue we came across related to the impact of market movements affecting collateral values, and hence exposure at default and loss given default. We felt that this is not a market risk but a credit risk, because it would only cause a loss if it was triggered by a credit event.

Outsourcing gave us some pause for thought. Generally this is considered an operational risk, but is the failure of an outsourcing counterparty a credit risk? We decided that if there are accruals, in other words if there are prepayments for services, then the loss in respect of these prepayments is a credit risk, but the wider implications of third-party failure should be considered under operational risk.

Linked to outsourcing is failure of an asset manager, including an open-ended Investment Company (OEIC) manager. We felt this should generally be treated as an outsourcing failure under operational risks except for reinsured fund links. We classed failure of a reinsured fund link provider as credit risk because of broader exposure issues in relation to reinsured funds – they are not ring-fenced in the same way as OEICs, for example.

Turning to insurance risk, we have 28 categories again, roughly based on Solvency II. Many of these are general insurance categories. We have broken them down into two different types of claims frequency and three different types of claim severity.

For claims frequency we have prospective claims in terms of the future claims emerging, and incurred but not reported (IBNR) in terms of past claims yet to be reported.

For claim severity, we have prospective claim severity in terms of future claims arising, the severity of claims that have been reported but are yet to be settled, and the severity of IBNR claims.

In terms of the demarcation issues, as I mentioned before, is non-disclosure a form of fraud? If so, you might class this under operational risk. However, it may be that non-disclosure is the fault of the life office in terms of poor wording of underwriting questions. We opted for non-disclosure as being an aspect of insurance risk on the basis that it will be implicit in your actual claims experience. I am sure many people might argue the point on this.

Another demarcation issue we identified was option take-up rates. Option take-up rates will vary with market conditions. The question is: are varying take-up rates a market risk or are they an insurance risk?

We took the view that ideally the assumption for option take-up rates should be dynamic. Any variance from assumed dynamic rates should then come under insurance risk.

Other grey areas we identified were expenses and, for general insurance, property rebuild costs. These would be linked to inflation, which we have classed as a market risk. However, we felt that there were factors other than general inflation that contribute to expense and rebuild cost inflation. Therefore, we have kept these within the insurance risk category.

In terms of liquidity risk, there is the question of what liquidity loss looks like. We identified that liquidity risk could give rise to assets being realised for less than balance sheet value as part of a “fire sale”. This is one aspect of liquidity risk loss.

The other aspect is interest on borrowing to tide over liquidity shortfalls. To what extent is this interest a loss rather than an adjustment for time value? We felt that only the excess of interest over base rates, or the risk-free curve, should be considered as a liquidity loss.

In terms of demarcation, there was a question regarding lines of credit. If somebody refuses to honour a line of credit and does not lend you money when they said they would, is that a credit risk?

Is it a liquidity risk? I think under a strict definition of credit risk as a failure to meet obligations, you could say it is a credit risk, but we have opted to treat this as a liquidity risk.

There are also demarcation issues regarding default of deposit counterparties. The first order impact would come under credit risk. However, there would be a second order impact on liquidity in terms of available funds that we would class as liquidity risk.

In terms of operational risk concepts, we started from the ABI Operational Risk Consortium (ORIC) definitions which are the same as the Basel II operational risk definitions. In terms of breaking these down, we considered whether we could use Basel II decision trees as a means of parsing individual risks. However, we found a few issues with these decision trees. For example, if you had a mis-selling event with a regulatory fine associated with it, the decision tree would put it into the “suitability, fiduciary and disclosure” category, rather than “mis-selling”, because there is a regulatory fine associated with it.

We felt the Basel II decision trees were not always suitable for consistent classification of operational risks. Instead we tried to identify as many operational risk categories as we could to help ensure consistent classification. It came to about 340 at the last count.

Moving on to strategy risk, I suspect this is not something actuaries would usually come across. We identified two different types of strategy risk. Firstly, there are exogenous factors, such as the Retail Distribution Review (RDR) or the actions of competitors, which are external to the firm.

Secondly, we have endogenous factors, such as the quality of our current product range and project failures such as failure to launch a new product, which are internal to the firm.

Within the endogenous factors we include brand and reputational risk. This includes not just the reputational risk of brand damage, but whether a brand actually supports the strategy.

Strategy risks do not just relate to new business, but also to back book initiatives. Say, for instance, that we have an initiative to achieve expense reductions, or to improve persistency under a customer loyalty programme. There is risk that that these initiatives fail to deliver, and anticipated benefits are not realised.

When we were discussing strategy risk, there was a strong view that there is no such thing. This argument maintained that strategy risk is merely a manifestation of other risks. For instance, the impact of equity market movements on sales is actually just a form of market risk. Similarly, the reputational impact from operational loss events, and the impact this has on sales, is an operational risk.

A related argument was that project risk, rather than being a sub-component of strategy risk, is a category in its own right.

In the end, we decided to retain strategy risk as a distinct category. Take the example of the impact of, say, equity market falls on ISA sales. You could say that this is market risk but, in practice, the controls you would have around the impact on new business would be completely different from the controls you would have around, say, the VIF impact of market falls.

We retained the strategy risk category but it remains a source of discussion between us.

Frictional risk, as I have mentioned before, relates to changes in regulatory capital rules and accounting rules, and the inefficiencies these can give rise to.

Finally, we have the aggregation and diversification category. As it stands, it is not a series of risks per se, but more a series of particular scenarios.

Generally, firms will allow for diversification across different asset types and across different risks. They will anticipate diversification benefits but these might not be realised.

To conclude, I would say there are many demarcation issues here. Looking at all the different issues that we came across, I would say there is probably no perfect system of classifying risks. I am sure there are plenty of other equally valid systems.

The one thing I would say in this system's favour is that it is a non-proprietary classification system freely available to all. I describe it as open source. Any actuary, as long as you acknowledge the Actuarial Profession's intellectual property rights, can feel free to use this. I am not saying you should use it, but when you talk about risks with peers, please bear in mind that it is unlikely that they will see completely eye-to-eye with you on risks and their terminology.

There will invariably be some subtle differences, particularly when you drill down further, in terms of risk classification between different organisations. To avoid possible confusion, in talking to peers from different organisations, I would ask that actuaries either use this common risk classification system or alternatively define in detail what you mean by a given risk term.

A final caveat is that risk classification is only a starting point in enterprise risk management. It helps in terms of allocating responsibilities and in reporting losses. However, you should always take care that you do not look at risks in isolation. There is considerable interaction between risks. You need to take a holistic approach to the entire portfolio risk rather than becoming too focused at the micro level.

With that, I will hand you back to the Chairman.

The Chairman: I should like to invite Mr Gordon Wood FFA to open the discussion.

Mr G. C. Wood, F.F.A. (opening the discussion): I think that there are two or three big picture questions that we need to ask ourselves. Does the world need a classification system at all? Do we need another one? Do we need one just for actuaries? If so, then is the proposed one the right one?

Many years ago, the company for which I was working was taken over by another life company which was owned by a bank. The bankers were not quite sure what to do with me so they called me "risk manager". My view at the time was that it was more important to identify all the risks, put controls round them, and measure them; and what risk buckets you threw them into for the purpose of aggregated reporting was of much less interest. I did come to realise, though, that risk identification is necessary but it is not sufficient, particularly as at the time I was working for a company owned by a bank. The bankers had a different risk language. The risk taxonomy was vital to get right. I did come round to thinking that a classification system would perhaps be useful.

Mr Kelliher touches in his paper on a number of existing risk systems. He mentioned one or two individual groups and also FSA and BaFin. I was slightly surprised that he did not talk about the

obvious one, the Committee of Sponsoring Organisations (COSO) internal control framework. This has been around since the mid-90s and was refreshed in 2003/2004.

There are quite a number of classification systems out there. What is special about our one? Why do we need a new one? In particular, Mr Kelliher starts with a few definitions. We have embedded value definitions, for example.

I might have liked to have seen a link to other industry definitions. With embedded value, for example, the European Chief Financial Officers (CFO) Forum has perfectly good definitions and high-level principles. We also mentioned the ORIC definition that the ABI use. We could simply start with that and then diverge to get the first appropriate version.

I should like to have seen more of the change control process. How do we ensure that we take other risk professionals with us? How do we ensure that our taxonomy adjusts and adapts as the ABI's or the CFO Forum's or adapts as the world adapts?

I might take issue with one or two areas in the paper, for example, the 340 subcategories in operational risk. One of the benefits of this mentioned was benchmarking across different companies. With 340 subcategories, I hate to think how long a period would be required to have credible data in most of those cells. I think there is a practical issue here.

Having said that, from experience, identifying subcategories at that level of detail is useful when conducting risk identification exercises and scenario planning in the corporate governance fora. It gets the management to think in detail about the types of events that might have an impact on their institutions.

Section 10.3 says that the defective strategy is particularly difficult to define or properly classify. There are many other examples in the paper. I do not necessarily think this is an issue so long as everybody agrees that these are the definitions and we will stick with them.

In section 13, which labels alternatives, I would have liked to have seen more about selling the eventual solutions. There is no form, for example, of SWOT analysis. I might have liked to have seen a little more of the thought processes and the reasons for the conclusions.

I struggled with the completeness of Appendix J (aggregation and diversification risks). Many of the other individual risk categories are pretty well complete. Sometimes the authors cheat by having a miscellaneous credit risk, for example. I notice they felt more confident with market risk, as there was not a miscellaneous category for this. The aggregation seemed to me a slightly random series of scenarios. It did not quite match the completeness and comprehensiveness of the rest of the paper.

My final comment is that if you go on Google, or any other search engine, and type in "ERM framework", there are three institutions mentioned as having such a framework: COSO, the American Casualty Actuarial Society, and the Institute and Faculty of Actuaries. Therefore, we are starting to have an external reputation for our thinking on ERM frameworks. This piece of work helps to advance and cement that reputation.

The Chairman: Thank you, Mr Wood. The discussion is open to the floor. Could we have the first comment or question?

Mr J. A. Porteous, F.F.A.: I wanted to make a brief comment. The paper is entitled “Common Risk Classification System for the Actuarial Profession”. I work in pensions and I do not see how this risk classification is going to work for me. You miss out many areas that are crucial to us, such as employer covenant, and benefit risk. I wondered how you could extend your work to bring in pensions so that we are all speaking a common language.

The Chairman: Thank you, Mr Porteous. Mr Kelliher, would you like to respond immediately?

Mr Kelliher (replying): A similar point was made at the London discussion. Despite 28 categories of credit risk, one thing we have not covered is the risk relating to employer covenant. For a pension scheme, obviously employer covenant is a huge credit risk exposure. It is an omission and, if we manage to develop this paper, it is something that we would seek to add.

In terms of the other aspects of risk classification, we have tried to think about pension schemes. If you were to drill down into the operational risk categories, you would see where we take the generic ORIC categories, 20 of them, and try to say how they could affect pension schemes. I accept this is not perfect; it is just the start. I take your point that more could be done on pension schemes and this risk classification system.

The Chairman: Thank you, Mr Kelliher. I think it is a useful comment, given that the European regulator has just started a consultation on what might be called Solvency II for pensions. This includes a holistic approach where the covenant of the employer is to be taken into account along with the assets that belong to the pension scheme. Measuring these is going to be a challenge to us as a profession. So thank you, Mr Porteous, for the question. Are there any more questions?

Mr J. E. Gill, F.F.A.: It is good to see the detail that the authors have gone into in terms of helping people with improving the quality and depth of risk assessment.

I have a couple of points: one is in relation to the central issue of loss data collection and how having an improved system of classification might eventually lead to better data collection. Those involved in risk management know that that is currently exceedingly fragmented and, in some places, non-existent. How can we build on a common classification system to improve the quality of data collection?

My second point is this. I have a fear that if we have a classification that is for the exclusive use of the Actuarial Profession, it will not last long. I would encourage ways of marketing this framework across the risk management universe. If we do not market this, it will have a short lifespan.

Mr A. M. Eastwood, F.F.A.: I endorse most of the words in the “Conclusion” slide. Like Mr Wood, I think that there needs to be more of a case made for having a different system from everybody else.

Reading the paper brought home to me the fact that it is important to understand how risks have been categorised and different categorisations are appropriate for different organisations. Even in the same organisation it may be appropriate to group different risks in different ways, depending on the purpose.

The example that springs to mind is credit risk and how to deal with fluctuations in the market price of credit instruments – effectively the market risk associated with fluctuations in the price of

credit risk. If the purpose is to work out the correlations between different risks, it might be helpful to bring into account the market risk associated with bonds with pure credit risk simply because the two will be closely correlated.

When considering statistics relating to past events and, in particular, statistics derived from external datasets, it is important to understand how those external statistics were built up and to use groupings that are consistent with the datasets concerned, or at least allow for any differences when you are trying to model a distribution of the effect of the different risks.

I think it is important, therefore, if we do agree that this is the taxonomy that should be used by the Actuarial Profession, to remember that it may be appropriate to group risks in different ways for different purposes.

Again, like Mr Wood, I question the value in having 340 different subcategories of risk.

I quite like the flowchart, although the way the questions are asked may end up with risks being mis-labelled. But this is acceptable so long as you are aware of it. When trying to build a database of operational risk records and categorise them in different ways, consistency is most important.

If you have just 340 different subgroups, I think you will inevitably end up with a degree of subjectivity and hence inconsistent allocation of different events to different risk types. It seemed to me, particularly when it came to operational risk as there are so many different types, that the flowchart was valuable when recording losses.

Mr P. H. Grace, F.F.A.: Mr Gill drew attention to the fact that actuaries are not the only people working in risk management. This paper is addressed to actuaries and is being discussed only in the United Kingdom. But there are actuaries in other parts of the world who are working in the same areas. I feel we should enter into dialogue with our counterparts in North America and Australia, who I suspect are also working on similar definitions.

Prof R. S. Clarkson, F.F.A.: The one point that left me exceptionally uncomfortable was the first bullet point under section 4: it is an event-based classification as opposed to cause-based classification.

Section 4.1 gives the explanation why. There could be multiple causes and therefore it might be too difficult. I think a lot more could be said about the Northern Rock disaster in the paper. To say it was liquidity-caused is too simplistic.

The Chief Executive of the FSA, at the Treasury Select Committee investigation into what went wrong, said something along the lines of, "We saw Northern Rock as having a low probability of getting into difficulty."

This was partly because their accounts seemed to suggest that their bad debt ratio was lower than the industry average. However, it appeared in the Financial Times, the Sunday Times and the Sunday Telegraph, that the deputy chief executive of Northern Rock hid tens of millions of bad loans. They did not appear in the accounts. He was prosecuted by the FSA and was fined £504,000.

Common sense questions whether, given the strategy of Northern Rock, you can believe the figures you see in the accounts.

How far do you go back in looking at causes? In a paper that I wrote for the Journal of Financial Compliance and Regulation, I mentioned a case that happened in 1772: the Bank of Eire. The parallels are unbelievably close. However, there are no textbooks around showing what happened in terms of bank failures.

By talking of event-based rather than cause-based, I think this paper tends to distract our available actuarial manpower in looking at difficult events.

The Chairman: Thank you, Professor Clarkson. Mr Kelliher, would you like to respond to any of the questions that have been raised so far?

Mr Kelliher (replying): Certainly. I think there have been many comments about the 340 operational risk categories. Is it too much detail?

Having dealt with operational risk in trying to assess operational risk capital, one of the things I found most frustrating was, given perhaps 20 different operational risk high level categories, the same event could be put under three different categories by three different people. I think one of the most frustrating things I have had to deal with was this lack of clarity, and the fact that people did not understand what each high level category should encompass.

It was only by drilling down, and spelling out in detail what that particular high level category covers, that we can achieve some consistency in classifications.

Mr Wood made the point that we are not going to get 340 cells of operational risk data. But, for me, the point is having 20 cells of consistent data, and for that we need to ensure that people understand on to what they need to map the data. That is why we went to that length of detail.

I would stress that we have also tried to cover multiple kinds of institutions. We did pick up on Professor Clarkson's point about the Northern Rock CEO. Included in the operational risk categories will be either deliberate falsification of accounts or inadvertent under-reserving. Either way, the categories should capture what happened for Northern Rock with its bad debts; and what I understand happened with HIH General Insurance, and possibly with Quinn Direct in terms of reserving.

The key point is that there are many risks, maybe not 340 for a particular organisation, but I would not be surprised if there is something in the region of 150–200. You need to have some means of mapping these individual risks on to more tractable groupings of risk. That is why we went to the level of detail that we did.

There was a point regarding COSO. I must admit that I did not become aware of the COSO system until quite a long way into this. I do not think you can readily access it – you need to pay £75. At the late stage I became aware of it, I decided not to pursue it.

I would note that one of the principles of the risk classification system is that it is freely available to anybody in the profession and to all risk professionals. I would like us to be able to use this when we talk to other risk professionals.

I hope we can spread the word about this risk classification system beyond the UK Actuarial Profession to the Americans in the Casualty Actuarial Society, and so on, and beyond actuaries to

the Institute of Risk Management and others. The Institute of Risk Management's deputy director, Richard Anderson, was at the equivalent meeting in Staple Inn last week and he was interested in this system.

The intention is to push this beyond the UK Actuarial Profession and make it freely available to anybody in the risk space. I hope it will be disseminated as widely as possible and that it will enhance our brand in ERM.

Mr K. A. Miller, F.F.A.: I also should like to thank the authors for their efforts, and I agree the need to define the terms that are used in this area. I often find myself talking at cross-purposes with the risk people in the bank I work for.

I agree with the comments that Mr Kelliher responded to, that we should enter into dialogue with overseas actuaries and probably other professionals as well. I accept that we are in a competitive position, as are some of these other risk professionals.

I have just a couple of detailed questions or comments. One is about the way in which risks related to embedded value were drawn out. I think that you start off with your assumptions, and then there is the risk that these assumptions are not borne out in practice. You seem to miss the issue of expenses, although it is touched on in the appendix. It appears in the chart but does not seem to be mentioned in the sections on insurance, demographic risk, and so on. I think you intend it to be part of the business risk, but it is not mentioned in the text.

The other one I was thinking of was tax risk – not corporation tax risk but the policy-holder tax assumption. You are probably going to bring it in later as a frictional risk in the chart.

Mr M. C. Ledlie, F.F.A.: It was clear from the paper that the focus in producing the classification was on insurance companies. I think that the authors need to be clear on the scope of the work that they are doing and the extent to which it could be applied to other entities. If it is going to be extended to other types of organisation, then there should be more thought about the types of risk that might apply to those entities.

I echo the comments made by Mr Wood earlier. Having governance and maintenance frameworks for this document is going to be important to make sure that it changes with time, and responds to debate and to input from other international associations, so that we can feel confident that this document will remain relevant over time and that people can actively use it because of the care and attention that it has received.

Finally, I have a comment on the frictional and diversification categories within the framework. These are not categories that I have ever seen used in a risk classification framework. I suggest that we should think carefully before applying these as separate primary risk categories rather than their being factored in through other types of risk within the classification.

Mr P. S. Carr, F.F.A.: First of all I should like to congratulate the authors on a very useful piece of work.

The comment I should like to make is that I do not find the use of the Linnaean terminology helpful. Linnaean terminology and cladistics work in biology because species can be classified uniquely

using the underlying genetics. I am not sure that you can do the same with risk classification. It may be more helpful to treat risk classification in a similar way to cause of death classifications. When people die, occasionally there is one clear cause of death. But for elderly people, there may be a number of causes of death and the examining doctor assigns one major cause on the certificate. It is not always clear what is the major underlying cause of death and the doctor must rely on expert judgement. Sometimes risk behaves in the same way. Organisations may succumb to risks, i.e. die. Moving away from Linnaean terminology might help us advance our thinking.

Prof Clarkson: Having said at the beginning of my comments that there was an Australian connection, I did not mention it in the end.

Over in Australia, an actuary, Adrian Fitzherbert, who had extensive practical experience of investment management, came to the conclusion that the education syllabus in Australia was too dominated by mathematical models. At Macquarie University, the teaching was modern portfolio theory, finance, economics, Black-Scholes option pricing formula, and so on. As a result of his campaigning, I understand that the Australian actuaries have now put in a course on economic history. You need to look back quite a distance at banking failures, or other business failures.

Sunspots have something like an 11 year cycle. It is not all that different from the banking cycle. But people tend to look at what happened in the recent past, not noticing that similar events have happened in the more distant past.

Really I am saying that there is an overall category of risk I would add: scientific risk, putting faith in the classification by separating risks in terms of modern portfolio theory, looking at the events not the causes. I would say that a lot more emphasis should be put on past history, as in Australia.

The Chairman: I will ask a question of my own, which is about the purpose for which one might use this classification. One use might be trying to control risks, perhaps monitoring them for reporting back to the board or evaluating the performance of a particular part of the company.

Mr Kelliher, do you want to say a few words about the uses to which this classification might be put? This is also an opportunity to respond to one or two other points that have been raised since you last spoke.

Mr Kelliher (replying): I think the primary purpose of the classification system is that we as a profession have some common reference point that we can use to talk about risks. It is not a huge aim – it is just trying to create a basic framework for discussion of risk. It does not aspire to much beyond that.

I suppose you could also use it in wider terms, maybe for looking at your risk register and asking, “Have we covered everything?” If you can find something that is not on your risk register, then obviously it could prove helpful. In terms of ORSA and its assessment of risks, it could help too.

To come back to Professor Clarkson’s point regarding the scientific risk, one of the categories that I have tried to bring out in respect of all the different types of risk is model risk. Whether it be market risk or credit risk or insurance risk, we do have model risk.

We have certain models which we use to value liabilities. They are built on certain assumptions. Without any change in the economic profile, a reappraisal of that model and of its inadequacies could lead us to change our reserving; and increased reserves would obviously have adverse consequences.

As actuaries we need to get into the habit of continually questioning models. In terms of risk, the key control you have against the risk that your model is imperfect is model governance, and with it the continuing questioning of whether the model is still valid.

I do not believe that there is any such thing as a perfect model. All models are flawed. It is just a question how flawed, and trying to identify what the weaknesses are.

Mr Ledlie mentioned the insurer's focus. We have tried to address other finance institutions. However, I must admit, there are the limitations of the working party itself. I have a life insurance background. The few other people involved were mostly from an insurance background. Necessarily, it will be coloured by our own limited experience, but we have tried to address other finance institutions.

For instance, under credit risk you will see many of the categories are bank credit categories that are not that relevant to life insurance.

Mr Miller mentioned expenses. I have assumed expenses fall under insurance risk. Working in life insurance, dealing with the FSA definition and various life office definitions of insurance risk, you become attuned to the idea that expense is an innate part of insurance risks.

I imagine for banks it might be different. Banks probably would not have an embedded value. The concept of variations in expenses and the impact on their value in force is probably alien. They would, I suspect, view expense risk just in terms of variations in next year's expenses relative to budget. I suspect that they have a different perspective on it to us. That is partly because of our broader look at VIF, as opposed to merely book value accounts.

Peter raised the point about cause of death analysis. It is an interesting idea. I suspect what you are trying to get at, Peter, is where an institution fails through a multiplicity of risk events occurring. It is something we have tried to pick up in the aggregation and diversification category.

An institution could maybe survive a downgrading, but if that was then to trigger a kind of margin call, maybe the liquidity strain might become too much. Coupled with various other factors, you could be facing failure. It is something that we have tried to pick up with the aggregation category. Often the whole is greater than the sum of the parts.

Mr R. Austin, F.F.A.: Under Solvency II, there are two approaches that companies can use. They can use the standard formula approach or the internal model approach.

Under the standard formula approach, the European Regulator has set out various categories of risks. Under the internal model approach it is the responsibility of the particular company as to which risks it wants to model. I would be interested to hear your thoughts, Mr Kelliher. In terms of the work you have done in risk categorisation, are there any categories of risks that could perhaps be added to the standard formula categories or has EIOPA/CEIOPS covered them all?

Also, from your experience of what companies are doing on internal models in Solvency II, again, are there any categories of risks that companies are not giving sufficient consideration to?

The Chairman: One of the concerns that was mentioned last week was that a box ticking mentality arises as a result of having many categories, and that people do not see the bigger risks.

On the other hand, one of the advantages seems to be that if you have a set of boxes to fill in, then the collection of data over a number of different organisations is much easier to do and more consistent. I think these two things seem to be in tension with each other.

Are there any more comments that people want to make?

Mr B. D. Wood, F.F.A.: The goal of the piece of work was to seek a common language. But to the extent that we are successful in achieving a common language, we achieve other things, too, because as soon as we start using a particular form of language, it defines the way we think about the problem. Therefore, it automatically defines the nature of the solutions that we generate.

I want to echo the importance that we should not choose a word just because it is convenient, and because everybody uses the same word. It is important to realise the specific word used will define subsequent actions that take place, in terms of managing risk and thinking about – and generating – solutions.

The Chairman: Thank you for all your comments. I would now like to call on Dr David Bowie, who has kindly agreed to pull together all these thoughts and to close our session.

Dr D. C. Bowie, F.F.A. (closing the discussion): I, too, would like to join with all the other contributors in thanking the authors for taking the time to write the paper, and Mr Kelliher particularly for presenting it here this evening.

It is difficult to describe a paper on the taxonomy of risks as exciting, but a paper that contains lists and jargon, and talks about events that go wrong will appeal to an actuary.

When I was reading the paper and trying to gather together themes and thoughts across the contributions made this evening, two areas came to mind; these also mesh neatly with the three questions that Mr Wood posed in his opening remarks.

The first area can be summed up in the question: how does this paper fit with the strategic steps that the profession is making in ERM?

The second area of debate relates more directly to Mr Wood's third question about the specifics in the paper.

I'm personally quite a fan of lists. Others in the audience also seem to be equally keen; indeed it sounded as if Mr Austin was trying to encourage an even bigger list of risks for companies to have to work through as part of their regulatory submissions. Mr Martin also commented that the use of lists enables us to work our way systematically across the landscape in which we happen to be working, and to make sure we do not leave out any big areas.

Mr Gill made the point that a common classification as proposed in this paper will facilitate any collection of data across a wider domain.

The downside of lists is as Mr Brian Wood pointed out: there is a potential danger that having a list might limit the debate to some extent. If you stick to the list and to the language in that list, that might be all that you consider and you do not look beyond those items, perhaps.

There was some comment from Mr Eastwood and Mr Gordon Wood on the number of categories. If the number of categories is too large, there is a danger of inconsistency; if it is too small, Mr Kelliher indicated that the lack of clarity will inhibit progress.

The general point that I took from several contributions, including those of Mr Carr, Professor Clarkson and Mr Eastwood, is that perhaps some of the most powerful areas of learning from this paper centre around the debates on the demarcation. The category in which you end up putting the risk is not important. Discussing what has caused you to allocate each risk to a particular category takes you closer towards risk management rather than just risk identification.

It is when you start having those discussions and debating the causes of the risks that you have to consider the appropriate risk appetites and governance arrangements.

Several contributors wondered if it was a useful exercise to invent a new language rather than learn about the existing languages used in risk management.

Mr Ledlie made the plea that any language or any set of lists should be maintained so as to be flexible enough to deal with the future, the ‘unknown unknowns’ that are the risks that will emerge in future.

A possible difficulty with being too flexible and too encompassing in any definitions is ending up with a definition so abstract that it does not mean much. It is somewhat akin to the urban myth that the word “white” does not exist in the Inuit languages; because “white” describes just about everything in their environment it in effect describes nothing. The plea is then that we do not try to be too general with the definitions, but that we are active in maintaining the lists in future.

Several of the comments made the point that the materiality of some of the risks which appear quite a long way down a classification system may be hugely important in some contexts but not others. The purpose of the measurement is important to bear in mind.

Turning to where this paper fits with the Actuarial Profession’s strategic ambitions in risk management, Professor Clarkson reminded us of the important role that actuaries have been playing in managing commercial enterprises with the use of risk measures since 1762.

In the last couple of decades, with RAMP (risk analysis and management for projects) and a few other initiatives, we have started exploring ERM more actively and we are looking to recast the work we do in an ERM mindset.

Enterprise risk management is a discipline (and quite a new one in any formal sense) in the wider world as well as in the actuarial world. There are not many university courses on ERM, for example. The Actuarial Profession has made some concrete steps towards embracing ERM with the inclusion of an ERM option in the education syllabus and the Chartered Enterprise Risk Actuary (CERA) qualification.

Several contributors wondered if this paper was ambitious enough or if it was unintentionally rather too constrained by its insurance world starting point. Mr Kelliher acknowledged that it would be helpful to build on the paper with expertise from other areas and the sense from this evening is that it would be worth supporting such extensions.

Messrs Grace and Gill encouraged us not only to market this paper across the actuarial community outside the UK, but perhaps even more widely.

The proof of this paper is going to be if the Actuarial Profession uses the language set out; and if we use these lists to make faster and more effective progress in building on our risk management heritage than any competitors.

Once again I thank Mr Kelliher and his co-authors for their paper. I hope that this paper will spark the success we wish for the Profession in embracing ERM.

The Chairman: Thank you very much, Dr Bowie. I should just like to give Mr Kelliher the opportunity to reply. He has already helped us by answering many of the questions as we went through. I am sure that there were a few others towards the end of the discussion, and also he might want to respond to some of the things that Dr Bowie has just said.

Mr Kelliher (replying): First, I'll come back to Mr Austin's point about whether we should add more risk categories in to the standard formula. You could argue that there are many gaps in terms of what the standard formula covers. We should always be aware that the whole point of Solvency II, and the whole point generally of solvency assessment, is that we have enough capital to cover our risks. You can tick all the boxes in terms of the standard formula but you would still need to address all the risks which affect our business in the ORSA, even if they are not in the standard formula.

One of the key things from the discussion is how to take this forward. The working party put a lot of effort into getting this far. I would like to see it going further. What we have here can help us explore risk topics on a consistent basis.

I certainly would look to spread this into the actuarial profession across the world and across the other risk professional bodies, like the Professional Risk Managers' International Association (PRMIA) and the Institute of Risk Management.

We should also seek volunteers from a wider base, from people who work in banking and pensions, to make sure that we are not missing anything.

I see this as a growing document. I do not think this is the end, by any means. We will always come across new risks. We can try to identify as many risks as we can, but there will always be new risks arising. There are certain limitations to our knowledge.

To quote Donald Rumsfeld, what we are dealing with in our models is the known unknowns; but there are unknown unknowns out there that we need to be aware of. Much of the time, the unknown unknowns should be known because they have occurred in the past. Sometimes it will be something completely new.

We need to have an evolving risk system because it will always change. As soon as anything in the business changes, new risks will appear.

I will finish with a call for volunteers. If anybody would like to help me take this to the next stage or even to pass the torch on at some stage, please feel free to contact me. We can, I hope, build on this.

The Chairman: Thank you, Mr Kelliher. I believe that there is to be an ERM guide to be produced by the Profession shortly. Maybe the classification will be a companion volume to that guide.

It just remains for me to express my own thanks, and, I am sure, the thanks of everybody here, to the authors, particularly to Mr Kelliher for not only being one of the authors but for making the presentation this evening. I should also like to thank Mr Wood and Dr Bowie in for opening and closing the discussion, and to those of you who participated in the discussion.