

---

# Model risk – daring to open up the black box by the model risk working party

---

## Abstract of the London Discussion

[Institute and Faculty of Actuaries, Sessional Research Event, London, 23 March 2015]

This abstract relates to the following paper: Aggarwal, A., Beck, M. B., Cann, M., Ford, T., Georgescu, D., Morjaria, N., Smith, A., Taylor, Y., Tsanakas, A., Witts, L., Ye, I. Model risk – daring to open up the black box. *British Actuarial Journal*, doi: 10.1017/S1357321715000276.

**The Chairman (Mr A. N. Hitchcox, F.I.A.):** Good evening, ladies and gentlemen. Welcome to tonight’s discussion on the paper: “Model risk – daring to open up the black box”. The chair of the working party and today’s speaker is Mr Nirav Morjaria. He is global head of economic capital and model risk management for HSBC’s insurance business. He has led the development and implementation of an advanced model risk management infrastructure and framework across its whole insurance business and he has embedded this into the business and into the risk management processes of the group. He gained a Masters in mathematics and was also awarded the Watson Wyatt prize for the top financial economics student on his way to passing his actuarial exams.

Mr Morjaria is going to present to us a summary of the paper and then ask some open questions to start the debate.

**Mr N. D. Morjaria, F.I.A.:** I am speaking on behalf of the IFoA’s model risk working party. I am going to cover the following agenda, which is in line with the order of the full paper. First, there will be an introduction to the working party and its objectives, and what we mean by model risk and what can go wrong. Then I am going to cover how model risk can be managed, with an introduction to the model risk management framework and how it could have mitigated some major real-life examples. I am then going to talk about governance and model cultures. What does good model governance look like and how does it take into account different attitudes to models? Then I am going to talk about model risk measurement, that is, how to quantify how wrong the model can be. Finally, I am going to run through some conclusions and our suggested areas for further investigation.

The working party kicked off just over a year ago and we have achieved a huge amount since then. The objectives of the working party are to raise industry awareness and understanding of model risk, to develop a framework for management and measurement of model risk and to foster good practice around the governance and control of models.

So, what do we mean by model risk? First, we need to start with what we mean by a model. We took the definitions provided by the Financial Reporting Council in TAS M and by the United States’ OCC and Federal Reserve Board in their supervisory guidance paper on model risk management. It is “any quantitative method, system, or approach that applies statistical, economic, financial or mathematical theories, techniques, and assumptions to process input data into quantitative estimates”.

It is quite a broad definition but we thought it was a sensible one to use.

The next logical question is whether we need this risk in the first place. Why do we need models? The environments that financial firms have to navigate, and the portfolios and strategies they employ, are complex. We need models to help us make the right decisions as human intuition and reasoning are not sufficient.

The point about making the right decisions is the key focus of what we have looked at as a working party.

Given that we need models, we then define model risk. Again, we have taken the regulatory definition from the US OCC and Federal Reserve Board as we thought that it was quite a clear and usable articulation of what is model risk. Model risk is the risk of adverse consequences from decisions based on incorrect or misused model outputs and reports.

There are two main causes of model risk: one is fundamental errors in the models; the other is the incorrect or inappropriate use of models.

Why should management be worried about this risk? It is probably fair to say that model risk is generally less mature and less well understood as a risk type compared to your typical market, credit, insurance and operational risks. However, with an increasing reliance on models due to ever-increasingly complex and competitive financial and non-financial operating environments, model risk is becoming a much more pertinent risk with substantial financial impacts.

We have listed a number of well-known examples of models behaving badly. There is a mix of financial and non-financial examples. There are plenty more examples in the report. One to signpost is the West Coast Main Line bid. The Department for Transport used models to assess rival bids for running a railway line. It had errors in its models and the cost to the UK taxpayer was over £50 million. It had to rerun its tender and there was damage to First Group's business plan.

Some other examples with which you might be familiar:

- The London Whale – JP Morgan. That led to \$6 billion in losses.
- The Millennium Bridge had to be closed for 18 months. Remedial works cost £5 million. It was because they did not realise that the resonant frequency of the bridge was the same as that of pedestrians.
- A key contributor to the 2008 financial crisis was an over-reliance on Gaussian copulas. Those were found to be inappropriate and did not correctly predict what would happen. Everyone knows what happened after that.

Many other examples are specifically actuarial. In summary, the financial impact of model risk can be substantial. It is important that we look at how we manage this risk. The question is: have companies' governance frameworks improved since these events? Certainly the regulators have been issuing a variety of guidance since the 2008 crisis. There have been many new regulations, such as Solvency II and Basel III, that aim to improve governance and controls around models. I imagine many of you have done plenty of work on Solvency II internal models, which are complex models that require onerous governance. However, the counter-argument is that many of these examples are since the financial crisis, so it is not certain that we have definitely addressed the issue.

What we need is a framework for managing model risk. How can model risk be managed? We looked at using established enterprise risk management frameworks as a starting point for the management of model risk: risk appetite, risk identification, risk assessment, risk monitoring and reporting, and risk mitigation. From that we developed a framework for managing model risks based on industry best practice, regulatory developments, and good practice risk management principles, and specifically consideration of the nuances of model risk as a risk type. We need a policy that sets out the roles of the board, the risk management function, model owners, the relevant governance authority and internal audit around risk management. We need standards around modelling – these are the more detailed instructions for model owners, standards to develop their models to cover data quality, model change control, model use, expert judgement, limitations and so on.

Model risk appetite is probably new to many companies. Effectively, it is the board setting its appetite for model risk. It needs to articulate its appetite for model risk in a risk appetite statement with limits. The limits will vary depending on the purpose for which the model is being used. But what we would generally find is that the appetite for errors and hence model risk for those models that are most business critical for the purposes of decision making and financial reporting, etc., have to be the lowest.

Some of the metrics that you could consider in a model risk appetite statement are, for example, the extent to which all models used have been identified and risk assessed, the extent to which models are compliant with standards, the number of high-risk rated models, cumulative amounts of model errors and number of model errors, etc.

We need to make sure that we have an inventory of the models that we have in place in the organisation, capturing the key management information (MI) about each model, and similarly for key model developments, which can also pose model risk.

The next step is materiality filtering. We need to identify to which of the models and model developments that we have identified as sufficiently material we should be applying a more robust and comprehensive model risk management framework.

Typically, you would have a quantitative assessment for model risk materiality, which might be in terms of measures like profit, reserves, capital, asset values, etc., underpinned by a qualitative assessment which would be made by the relevant model owner. This would consider things like the extent of reliance on the results of the model for decision making, and how important the decisions are that are being made, etc.

We then considered quantitative model risk assessment and qualitative model risk assessment. Quantitative model risk assessment we will cover later, so I will not go into detail here. Qualitative model risk assessment is effectively an RAG rating, depending on the risk of the model.

Under risk monitoring and reporting, we need to make sure that we are developing model risk MI. We need to make sure that, as the board has ultimate responsibility for managing the model risk of the organisation, the board, or its relevant delegated committee structure, sees the model risk MI and is able to carry out effective oversight of model risk.

Finally, model risk mitigation consists of actions that we can consider taking if we are outside our model risk appetite, such as additional validation required. Or, it might be that we say we

are not comfortable with the model so we need to apply an overlay of expert judgement, or something else.

Let us consider how to apply this framework to a few case studies to try to bring it to life. We considered three case studies. The first one was the West Coast Main Line Rail franchise. Here is a summary of what went wrong. Basically, Virgin had been operating the West Coast Main Line up until the point of the franchise renewal in 2012 and both Virgin and First Group bid for the franchise. After assessments by the Department of Transport, the franchise was initially awarded to First Group. However, Virgin requested a judicial review of the decision because they had a number of questions around the forecasting risk models used by the Department of Transport. The Transport Secretary agreed that an independent review could be performed, which identified a number of errors. First of all, there were mistakes in the way in which inflation and passenger numbers were taken into account, and the values of those variables were understated by up to 50%. Second, the Department of Transport provided guidance to bidders but used different assumptions to different bidders, which created inconsistencies and confusion. Third, economic assumptions were checked only at a later stage when the Department of Transport calculated the size of risk bond to be paid to the bidder, and there were various errors found in those risk evaluations such that the risk was understated.

What went wrong in our model risk management framework? There were a number of things. We mentioned data quality and expert judgement. Mistakes were made in the inflation assumptions. They should have been checked for reasonableness, particularly as the Virgin bid would have been based on data that was complete and appropriate. That resulted in clear bias to First Group, as it had assumed more optimistic values than Virgin.

The extent of model validation performed was inadequate as it did not identify the technical flaws in modelling or inconsistent assumptions. Fundamentally, First Group's model results did not sound reasonable.

Had all bidders used the same models rather than each using their own different models, this would have permitted the Department of Transport to compare the rival bids and identify any inconsistencies.

Finally, in terms of model governance, there was a lack of understanding of the model by the stakeholders, and overall insufficient transparency, governance and oversight of the project. Roles and responsibilities for functions, committees and boards within the Department of Transport were not clear, and there were early warning signs that things were going wrong when external advisers spotted mistakes. They were not communicated or formally escalated. Incorrect reports were circulated to decision-makers. Ultimately, there was a lack of accountability which resulted in the taxpayer paying the burden.

I am just going to touch on a couple of other examples. Many of you will be familiar with Long-Term Capital Management, which was a hedge fund. It had a model that it used and it did spectacularly well in the initial years. They achieved fantastic returns of 43% in 1995 and 41% in 1996. As the hedge fund grew, it decided to move into different riskier opportunities, such as merger arbitrage. But its models were not developed for those purposes. As a result of the economic crisis, many of the banks and pension funds that were invested in Long-Term Capital Management moved close to collapse, and in under a year the fund lost \$4.4 billion of its \$4.7 billion in capital through

market losses and forced liquidations. So, again, what went wrong in the model risk management framework? It was not just model risk, but model risk was a key contributor. There was an over-reliance on the model and the assumption that it would always work in the same way, even when Long-Term Capital Management moved away from its original investment strategy. In terms of model methodology and expert judgement, there was no transparency externally of either the investment strategy or the underlying model. Finally, in relation to model validation and external triggers, there was an inadequate review of the model and its objectives when the trading strategy was changed.

The third example is JP Morgan and the London Whale. JP Morgan made trading losses of £6 billion and they were fined £1 billion. What happened here? The chief investment officer was responsible for investing excess bank deposits in a low-risk manner. In order to manage the bank's risk, the CIO bought credit default swap derivatives which were designed to hedge against a big downturn in the economy. The bank increased more than tenfold its portfolio from \$4 billion in 2010 to \$51 billion in 2011, and then it tripled to \$157 billion in early 2012.

However, while JP Morgan initially intended to use the synthetic derivative portfolio as a risk management tool, instead it became a speculative tool and a source of profit for the bank as the financial crisis ended in 2011. Existing risk controls were flagged. In fact, the trades were effectively ten times more risky than the agreed guidelines. The CIO breached five of the critical risk controls more than 330 times. However, instead of scaling back the risk, JP Morgan changed the way it measured its risk by changing its value-at-risk metric in January 2012. So, not only did it change the model, but it made an error in how it did so. The error in the spreadsheet used to measure the revised value-at-risk meant that the risk was understated by 50%. In early 2012, the European sovereign debt crisis took hold, markets reversed and that led to trading losses totalling more than \$6 billion.

What went wrong in the model risk management framework? It was not just the model but that was a key contributor. The CIO was supposed to be investing deposits in a low-risk manner but the bank's activities meant that it exceeded value-at-risk limits for 4 consecutive days. Management claimed not to understand the risk model. They said it was very complicated, but a number of CIO staff were on the executive committee. In terms of model validation, when existing controls highlighted the level of risk, the bank just changed the model and subsequent investigations found that the spreadsheet used was not queried and not properly validated.

Finally, the internal JP Morgan report stated that CIO judgement, execution and escalation in Q1 2012 were poor and that CIO oversight and controls did not evolve with the increased risk in the increasing complexity of its activities. Subsequent Senate reports said that the CIO risk dashboard was "flashing red and sounding alarms". The regulator also came in for criticism for its oversight. This inappropriate activity by JP Morgan was first noticed in 2008, but was not followed up.

That was a number of examples. I hope the model risk management framework that we suggested will have helped to improve the situation. The model risk management framework can only ever be as good as the people who use it, and that will depend on the model risk governance and culture of the organisation.

So, what does good model risk governance look like? Governance and controls to manage model risk should build on existing practices rather than replace them. But we think they should extend

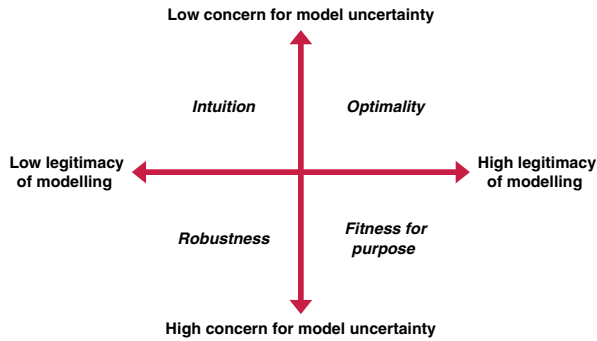


Figure 1. Model perceptions

beyond the regulatory perspective. Current governance and control levels around models include the separation of “do” and “review”. We also have specifications for requirements, technical design and testing, reconciliations of data, oversight of assumptions and sign-off and control of models and so on.

Solvency II and similar regulations are looking at enhanced governance controls including formal recognition of model risk, so that it is managed like other types of risk. Also, there will be established governance frameworks for model development, implementation, use and validation, subject to internal audit scrutiny. We have started the idea of model inventories and documentation standards and making model limitations transparent and well-articulated. We have things like the use test and independent validation. That is where the industry is going.

So, what does even better model risk governance look like? We think it looks like building on the above but with culturally aware governance and controls. Governance should require people with different perspectives on model risk to all contribute as part of the decision-making process. There needs to be two-way challenge that is supported and accepted. We need to consider proportionality, so the materiality of the decision will trigger the extent of the governance activity.

In order to understand the ways in which model risk arises, we need to consider the different ways in which models are used in decision making. Our proposed categorisation is seen in Figure 1. The horizontal axis reflects the perceived legitimacy of modelling. In the right half plane it is believed that models should be used in decision making such that the emphasis is on measurement and computation.

In the left half plane models are given an insubstantial role in decision making. Here the emphasis is on intuition and expert knowledge. The vertical axis reflects concern with uncertainty. Stakeholders in the top half plane are confident in their processes leading to good decisions and are generally not concerned about model uncertainty. Agents in the bottom half plane are less sure of themselves. In particular, model uncertainty is a major concern for them.

Each quadrant gives a different way of using models. At the top right, agents are keen to use those sophisticated models in order to take decisions that are optimal. For them, decision making is a process that can and should be driven by modelling. At the bottom right people are primarily concerned with the technical validity and fitness for purpose of the model in use, possibly a more actuarial perspective.




Figure 2. Risks associated with model perceptions

At the bottom left, stakeholders view all risks that matter as ever-changing and essentially impossible to measure; thus they want decisions that are robust to such uncertainty. For them, models can never really be fit for purpose and they show a preference for stress and scenario testing.

Finally, at the top left we have decision-makers who do not see why models should be used in the first place. For them, intuition and deep knowledge of market realities will always trump mathematics. Modelling is just an expensive distraction.

Each different way of using models – or indeed not using models – brings with it its own risks (see Figure 2). In particular, if one of the four perspectives dominates within the organisation, the following might happen. At the top right the danger is of excessive reliance on models with little concern for their limitations. Evidence that is inconsistent with the models may be disregarded for too long. At the bottom right there is a risk that too many constraints are placed on the use of models to derive business benefit. A different risk is that the overall modelling framework is flawed.

At the bottom left, decisions that are robust to a high degree of uncertainty will generally be suboptimal. This can lead to an excessively conservative way of running the business while opportunities identified with the help of models are missed. At the top left the risk arises from over-reliance on intuition and market knowledge while ignoring the information and insight that a model can bring. Decision-makers in this quadrant, if left unchecked, may provide incentives for developing models that just generate convenient outputs. In this scenario, the major risk is loss of accountability. If intuition fails, will it be recognised as such or will models or modellers take the blame?

We believe that each of these four perspectives brings something useful to the table, but also that each, if left to dominate, produces substantial risks. When faced with deep uncertainties, all four perspectives must be taken into account and respond to each other to have successful model risk governance.

So, finally, we are going to touch on how to quantify model risk, that is, model risk measurement. We discussed the idea of quantifying the model risk inherent in a model or model development being challenging and a relatively new concept. However, it is crucial in effectively managing model risk that an organisation should understand how large a financial loss its models could cause.

Models can contain several types of errors ranging from statistical uncertainty to human blunders. Processes exist for understanding the impact of these uncertainties on model output.

To quantify model risks in the context of a firm's risk management, we need to assess the financial impact of model errors on a firm's profitability. Model output may contain a large error, but the losses to the firm may be small if the model output is ignored or if management are aware of model limitations and adopt strategies to mitigate model risk. On the other hand, a small model error may have large consequences if important decisions rely on model output.

So we might assume that expected model errors are 0 but human blunders are as likely to lead to overstatement or understatement, and that biased statistical estimates are correct on average. The expected model error may be 0 if, for example, a model is as likely to understate as to overstate the output and if corporate profits respond linearly and symmetrically to model mis-statement. On the other hand, there are costs associated with model changes such that the expected financial impact may be related to model output volatility rather than the expected value.

In general, we may then want to split model error into two components, so, as is typically the case with credit risk, we consider expected and unexpected risk, so expected model error is the losses expected under a model containing a baseline degree of model error. Unexpected model error is the contingent financial impact of higher than expected model errors.

Financial impacts of model errors are generally asymmetric so one way to quantify the model risk might be to use option pricing techniques.

We considered some specific "deep dives" of types of models. We tried to look at how we would quantify the model risk for those models. Proxy models are unique in that the model error can be precisely quantified, and we would use rigorous out-of-sample testing in order to be able to do so. We would monitor the model error for each business line against stated appetite limits and look at whether we remained within those limits.

In longevity models, the model risk arises for two reasons: first, if the chosen model is found to be invalid or, second, where an event occurs that is outside the historical data set. We considered stress and scenario testing of models and events that are not in the data in order to provide a reasonable range of outcomes.

With investment illustration models, we can carry out Monte Carlo back-testing as extensive historic asset return data is widely available. We can supplement that with stress and scenario testing again to compare against stated confidence levels of meeting investment aims – that is, our model risk appetite.

Finally, we considered environmental models, which are a little different but we thought it would be useful to consider non-financial models. Interestingly, there are a number of similarities to financial models in the model risks to which they are exposed, and in how to measure and manage these model risks.

I am just going to sum up with some key conclusions. Model risk results from model errors or approximations or from inappropriate model use. The impacts, as we have seen, can be substantial.



Model risk can be managed through a model risk management framework. We have tried to define what good practice could look like. But, importantly, that requires board engagement and support to set an appetite for model risk and to manage model risk within those limits.

We then talked about culturally aware governance and controls as good practice. Optimal model risk management should consider all four of the perspectives equally, not just focussing on technical fitness for purpose, which has probably been the focus of regulatory regimes to date.

Finally, we tried to quantify the financial impact of model risk. That is dependent on how the model outputs drive decisions. There are specific methods that we could use to quantify model risks, which will vary for the different model types. Again, we would monitor the results and manage them against our risk appetite limits.

To close, we have set out a number of questions that might be useful to consider. We are interested in your thoughts as to what keeps you awake at night about your models and what we as an industry could do to address those concerns.

We have defined the framework for model risk management, building on existing ERM frameworks. We are interested in your thoughts on whether model risk should be managed in the same way as other risks. We would be interested to know if you capture model risk in your operational risk frameworks already.

Around the governance processes, how can we ensure that we capture effective challenge from all the four different quadrants? How do we ensure that all the perspectives are heard and responded to without ending up with just another box ticking exercise?

In terms of the areas for further investigation, we are going to look at developing the case studies that we have identified to illustrate how the concepts described in the sessional paper would work in practice. We are going to look at further developing the practical model risk governance responses in line with the four quadrants that we have discussed. We are potentially going to consider systemic risk. We have not yet considered systemic risk associated with the use of models, looking at what peers do to drive modelling practice.

Finally, we will carry out more technically focussed research associated with model risk measurement. We looked at, at a high level, how to quantify model risk but we need to do further work in order to be able to do so more robustly.

**The Chairman:** I am shortly going to open the discussion to the floor. Before we start the discussion, I am going to explain why I have a bottle of Champagne on the table. We all understand that model risk is important. What I find is when I use the term “model risk” to an underwriter or a director they say, “Oh, you mean the actuaries have got it wrong again”. That is unfair, and we have to do something about it. So, I am offering a bottle of Champagne to the best alternative sound bite or phrase to “model risk” that is snappier, that a non-modelling person will gravitate towards and that does not sound defensive.

**Mr N. Shah, F.I.A.:** My one point is that within your presentation there seems to be a mixture of basic operational types of error in model use, such as using the model for the wrong purpose, whereas is not the biggest model risk the mis-specification? That is, the model is just a bad

approximation of the world. I did not see you bringing that aspect out. I would be interested in your thoughts.

**Mr Morjaria:** I have a slightly controversial view. I would also be keen to get the views of others from the working party. I think that is the actuary's view of the biggest risk. We do consider a number of examples in the report of model errors, and those have caused significant issues. For example, the West Coast Main Line example we talked about was of exactly that type. It was an error in the coding of the model.

We started from that premise. We looked at the chronology of many different case studies. We found that while there have been many such incidents, a large number of the incidents have been related to the use and application of models. There is often a breakdown in communications between different parties. Often the model users do not take sufficient responsibility for the models. They rely on the modellers, the actuaries, and they have ended up using models for purposes for which they were not originally supposed to be used or not properly considering limitations, and so on. There have also been instances, such as the JP Morgan example, where the CIO was always aware that he was breaching risk limits and so he decided to change the model because he could. There was not enough governance or validation around changes to the model.

I agree with the point that the core risk is around errors in the model. That is why we have defined model risk as errors in the model, as well as inappropriate or incorrect use of the model.

Are there any additional thoughts from the working party members?

**Dr A. Tsanakas:** While I agree that mis-specification is a key risk with models, I would also agree with Mr Morjaria that this is a particularly actuarial view of framing the problem. Talking about mis-specification of models already makes implicit assumptions: that there is some sort of discoverable risk process to which an insurance company is exposed, such that nature throws its dice in some orderly way. Then we just have to wait long enough and do all our jobs properly to discover what are the real odds (the "true model"). This is a valid interpretation, but it is not the only way of viewing uncertainty.

For example, if you consider the possibility that the risk environment changes over time in radical ways, then you have no hope that the model will ever be sufficiently well-specified, which is exactly why you need a diversity of perspectives, as opposed to just the focus on the specification, in order to make sense of the nature of the problem itself.

**Mr Shah:** I think my comment really related to the coding, and so on. Those risks tend to be picked up in operational risk. I was clarifying to make sure we do not double-count these risks. But I do understand why you put them where they are.

**Mr Morjaria:** There is a risk of double-counting between model risk and operational risk. I am not saying that we are right but we think a subset of model risk falls into the operational risk category. There are certainly nuances that we would not necessarily capture through operational risk processes. So, yes, I think that was a valid point.

**Dr L. M. Pryor, F.I.A.:** I think the paper's key contribution is the introduction of actuaries to these cultural theories of thinking in a systematic way about the different ways that people think about models.

As the authors point out, we, as actuaries, possibly tend to drift towards the bottom and right, and I suspect that many of us are about as far bottom and right as we can be in Figure 1.

I think the other particularly valuable part of the paper was section 4.5, comparing the way that insurers in particular, and financial institutions in general, look at and use financial models, and the way that environmental models are used in slightly different contexts. The financial and environmental models have much in common. They are typically complex and cover the long term. Also, the users and other stakeholders of the models have a lot bound up in the outcomes, emotionally as well as more rationally and economically. I think that drawing these parallels and learning from the ways in which people have thought about environmental models is useful for actuaries.

I also liked the general fitness for purpose theme pervading the paper with the emphasis on what is the purpose.

Returning to the bottom right quadrant, I read some of the comments on how people in the bottom right quadrant concentrate on fitness for purpose without worrying too much about what is the purpose.

To my mind, it is the purpose that is important. The fitness bit in some cases is just tinkering at the edges. You might have slightly better parameterisation and slightly better validation; but if the whole concept of the model is wrong for the purpose for which you are using it, then you might just as well shut up shop and go home.

I was worried about the emphasis on the quantification of model risk. You are in danger of falling into an infinite regression in that the models you use for the quantification of model risk are themselves subject to model risk.

The most positive outcome I think this paper could have would be to encourage actuaries in general to move outside their comfort zone. Quantification is often comfortable for us. Let us try to move outside that zone sometimes.

On that note, I should just like to answer one of the questions raised: What keeps me awake at night. This is not necessarily about the models that I build, but is about models in the financial services sector in general. It is the big changes from external triggers that we have not thought about. What about people's sales projections for annuity business now? How are they different from the ones that they were making 13 months ago? What about all those banking risk models – the foreign exchange models? How are they different now that the Swiss franc is no longer pegged?

It is those sorts of changes that are going to have huge effects. The regulatory changes are not necessarily economic or financial assumptions; but they are assumptions that are vital for the usefulness of the models. So I think that is what would keep me awake at night.

**Mr A. Poracchia, F.I.A.:** I have a question for the working party. To what extent have you considered the time dimension to the definition of risk? What quantifies it for simple models, such as those for reserving? A model error can go unnoticed for quite some time before you realise that there is one. So what is the definition of risk in terms of measurement? I think that would link back to some of the potential double-counting issues that were raised earlier.

**Mr Morjaria:** This is a difficult question. We probably have not thought about it sufficiently. It is something we will take away to consider as a working party.

**Mr H. T. Medlam, F.I.A.:** I have a question for the working party. Would you rather have a simple model that you can understand and know its limitations, or a complex model that better models the data?

**Mr Morjaria:** I think the answer to that probably depends on where you sit in the model risk culture diagram. Personally, I prefer the value in simple models. They are easier for users to understand.

**Dr Tsanakas:** My inclination would be the same as Mr Morjaria's: to go for simple models that we understand. But we see, in practice, that different stakeholders prefer more complex models. This may be because such models can give higher granularity, more MI or more detail to satisfy reporting requirements. This is, in some sense, part of the negotiation that is taking place. You may want something simple so that you can better control the error. Somebody else may want something more detailed so that they can have the output that they need.

**Mr R. Kelsey, F.F.A.:** I did some back-testing on the examples. For Long-Term Capital Management I also thought of Nick Leeson, who did not have a model but he did quite a good job at bankrupting a bank without it. I do not think that it is fair to concentrate on model risk. In most of these cases, controls have been overridden by people.

If I were asked to try to win a bottle of Champagne, I think I might offer "obsolescence risk". Models go through a development curve. To begin with, they are not trusted. You do a certain amount of testing, gaining an increased amount of confidence with the model. After some time, you become blasé. You might update it and make a mistake or leave it. But just like the underwriter who made some pricing mistakes, or Nick Leeson, I think you become obsolescent.

**The Chairman:** And if I can put words in your mouth, once you have used the model for 5 years, you tend to become comfortable with the hidden assumptions that you made a long time ago, do you not? That is an interesting observation.

**Mr Morjaria:** One specific point: that model users changing their models does constitute model risk. What we are saying is that there is not sufficient governance in the process of changing the model purpose. I acknowledge what you are saying, however. It is certainly not just about the model itself.

**The Chairman:** When I looked at my senior management team and my board, I could put named individuals in each of those boxes in Figure 1. Obviously, I am in the bottom right-hand corner. But I can put my underwriter, my chief operating officer and my chief financial officer all in different boxes. It fortunately turns out that my company has an even spread across all those characteristics.

But then the penny dropped – it is my job as Chief Risk Officer to make sure we have an even spread. If you are the CRO, and you are assessing your board's ability to understand the model for Solvency II purposes, you have probably done diagrams, you have a liability underwriter and a reserving actuary, you have a finance person and you have an operations person.

But I think you should also look at your board members and make sure that you are not too strong or too weak in one of those boxes.

**Mr M. Thompson:** I am to blame for this fourfold diagram. At the board level you have all the boxes nicely covered. I have done some work with Mr David Ingram, who has been going into insurance companies and doing interviews. He has spoken to the board and to the next level down and then gone down again to middle management.

It is not the same pattern at each level. There is a nice even balance at the board level, perhaps, and then it is quite out of balance, by comparison, down at the next level. When you reach middle management, it changes again. In some of the companies that he looked at it was the bottom left quadrant, the one where you are extremely cautious, that was dominant.

When he presented those results to the company, they said, “That is what has been going wrong. Now we understand the problems we are having”.

**Mr P. Kaye:** For the Champagne, I am going to offer “model culture risk”. From what I hear, it is about balance, and even the word “legitimacy” here seems loaded. It says you have to be more reliant on models.

The question I have for the working party is about the idea of measuring models. Sometimes, of course, you should never have modelled it in the first place. It is almost misleading to have done so. What about promoting the issue of looking at model resilience? If all the models were nonsense, how would you cope?

**Mr Morjaria:** Just to understand that better, are you saying that if we think about whether all models are fulfilling their purpose or not, that would be a useful avenue to explore effectively?

**Mr A. D. Smith:** First, let me congratulate you on identifying the elephant in the room, which is that you can use models to describe other risks. If you use models to describe model risks, then it is circular and you can go round that iteration many times. You definitely have not captured all of those risks.

We have done some thinking about the model robustness question. There is something you can do that is quite helpful in this regard. You can say “Here is a process that I am following to model this particular phenomenon; but I could be completely wrong about what is driving it and it could be some alternative reference model. So I need to go through a list of possible explanations of reference models and then see how what I am doing responds in each of those cases. If any one of them is disastrously wrong, then I have to modify the way I am looking at it”.

That is turning on its head what you might think of as the more traditional actuarial perspective of working out what the right model is and then using the right model. You are never allowing yourself to focus on just one model. You are always saying there are many things that could be going on here. I want to make sure that, whichever of these is going on, what I am doing is either sensible in the first place or at least can respond to me having been wrong.

**Ms K. Gilewicz, F.I.A.:** We talk about calculating and capturing model risk. But for me one of the key issues is communication of the uncertainty in our model. I would be quite keen on hearing views from the working party on this subject.

**Mr Morjaria:** We have identified in the paper a number of possible questions that we could be asking or could be considering in terms of the information that we would want to communicate relating to uncertainty. I agree that that is an important point.

**Ms C. Lam, F.I.A.:** When building a model the modellers always have a set of beliefs. It is based on those beliefs that they build and develop the models. How often do the modellers go back and check those beliefs correspond to the interactions in reality?

I have built models and that question is always at the back of my mind. Once you have built a model and do the validation, you do not always go back to the very first point to check that. Are there any practitioners here who can shed light on this topic?

**Mr Shah:** I think that it goes back to one of the comments made about model obscurity. I think that the whole modelling process is a question of trying to find a simplification of the real world, and as you learn more and more, the model evolves. I have been using models for quite a while and it is a continuous process. There are times when you figure out a new modelling approach that suits the current work, but the old one suited the position as you saw it at that time.

**Mr M. G. White, F.I.A.:** Responding to Mr Smith's point, it involves accepting what you do not know, which I think is a positive step. But it is worth remembering that what users and clients, and even the public and political leaders, are often asking is: "You are the expert, can you not tell us the answer?". It requires courage to respond appropriately.

**Mrs E. J. Nicholson, F.I.A.:** I wanted to ask about the extent to which you have thought about how validation is successful in mitigating risks, or the extent to which validation omits a portion of that risk or overlooks it and leaves you with a residual model risk of which we should all be more aware.

**Mr Smith:** The validation process is, I think, most helpful for uncovering either human blunders or deliberate dishonesty – somebody who has falsified the data or has made an unjustified choice of data or has just messed up so that the formula has a multiplier where it should have a divide.

Where the validation process seems often to miss things is where there are many potential models that could explain the data. Too often the validation process involves taking the model you picked and doing a statistical test against the data and then not rejecting the model while being blind to the possibility that there are 100 other models that should also not be rejected.

We too easily confuse the notion of "not rejected" with "proven true". Probably the word "validation" is slightly unhelpful and does not always describe clearly what it does. With a finite amount of data, on those terms you cannot fix that problem. There are 100 different models which could explain those data points and you are never going to be able to eliminate 99 of them. Perhaps you have to wait for 500 years' more data to be able to eliminate 99 out of 100 of them.

There is certainly a concern that the word "validation" could give a false sense of security as to whether a model is right rather than: "Here is a list of avoidable pitfalls and we have managed to mitigate some of those".

**Dr Pryor:** Returning to what Mr Kaye said earlier, thinking about how we can be sure that the models are correct, what about reverse stress testing? What would have to be true in order for a particular model to be useless, and not letting yourself be swayed by arguments such as, "But that could never happen"? That is difficult.

So, what could never happen? In other words, what could happen that would stop the model being useful or stop it giving useful results?

My shot at the Champagne is “terrible trust in technology”.

**Mr Shah:** Just coming back to Mr Smith’s point, I think traditionally validation has been seen too narrowly as the justification of a model. It goes back to whom you should involve in the validation process, and that chart again. Diversity is the key to validating a model. Effectively, all the other possible models are looked at by a diverse set of people who can challenge more creatively whether it is directionally correct.

**Mr N. E. Sheridan, F.I.A.:** I would like to obtain your opinion on this point. I think there could be some value in further narrowing down what is a model. The definition we heard at the beginning included everything from adding two numbers together upwards. Some sort of further categorisation could be useful.

Also, looking at the list of failures, some of them you looked at were where model review and better governance could have been helpful but there were others where the models used were perfectly legitimate. For example, the Gaussian copula models were well understood and well validated at the time, but we could not foresee the future despite this fact.

Is there also some sort of categorisation of those failures where some originated from model risk and some originated from not being able to foresee the future?

**Mr Morjaria:** The definition of a model is subjective. The definition that was taken is probably the purest regulatory definition.

The way that we looked at narrowing things down was more through the application of the materiality filter so that, even if you start from an initially large list of what would come under the definition of models, the intellect comes in applying the materiality filter effectively.

**Mr T. G. Ford, F.I.A.:** I am a member of the working party. I was interested in the comment made that we had the breadth of the model definition set widely. One of the concerns that we had was that the potential for model risk is exacerbated in those areas of a business that might be subject to less scrutiny. So, if you look within the finance or actuarial areas of an insurer, where processes are subject to both internal and external audit, we felt many of the problems with models are already well explored. Also, internal models under Solvency II receive much attention and the necessary focus is already given.

However, if we look at other models, such as those that are operating on the insurer’s website – including financial projection models looking at affordability of retirement – less scrutiny is currently given. So, our emphasis was more on keeping sufficient breadth so that we did not potentially miss those models subject to less focus where there is a lack of appreciation that model risks might exist.

Finally, I think that the core part of the definition of a model is uncertainty and the use of assumptions, although this is not as broad as covering all the spreadsheets within the organisation.

**Ms M. J. Strudwick, F.I.A.:** I just wondered if the catchphrase that we are looking for is “get real”. Where I have seen validation to be most effective is where there is a discussion with people who

understand the business, understand the different modelling approaches, and challenge whether the model and assumptions – implicit or explicit – that have been selected are representative of the real world. I think where that is done well it can be valuable.

I think it is a useful paper. I wonder whether we are limiting ourselves, or the working party is limiting itself, by just focussing on financial models in terms of its audience, and whether this paper might achieve wider reference throughout other industries.

**Mr Morjaria:** I like the suggestion of “get real”. That gets to the point. The point about wider models is a good one. We have tried with environmental models. Initially, when we were discussing the topic of model risk, we were cognisant that it could be applied more broadly than to just financial or actuarial models. Given our bandwidth and expertise, we were initially coming at this from the financial perspective.

I agree that if we look at a second phase, then we would want to consider broader models.

**Mr Medlam:** I studied engineering at university. Engineering modelling is different in that you are modelling event outcomes and the events are repeated time and time again. Engineering models are calibrated to a fixed frequency or an operating point, as it is called. The models work well around that operating point, but as soon as you move away from a certain operating point, the models breakdown.

Going back to Dr Pryor’s point of working out where the models breakdown, understanding the operating point and operating range of our models is important for the business. We need to be aware of and communicate when the models do not work as well as when they do.

**The Chairman:** Let me throw a provocative comment at the audience. Model risk is big, is it not? You cannot open a paper these days without finding somebody who has made a mistake with their model costing billions or hundreds of millions.

So the question arises about disclosure of your model risk. Those of you who read the FRC’s documents will know that all stock market companies have to disclose their principal risks and uncertainties starting from next year-end. So it seems to me that if you have a model behind your embedded value in a life insurance company, then one of your biggest risks to disclose is your model risk or the extent of it.

What do you think?

**Mr Poracchia:** This is going back to what is a model. We carry a voice in industry and everything actuaries do is about modelling in a way. There is now a different generation of models: those that are hypothesis based, trying to model what you believe, and then there are many new techniques emerging that are more statistical based – to learn from the data rather than make assumptions. I think that we have to be careful when we start putting forward a few points of view as to the size of the risks and what that will bring back to us as liabilities and responsibilities. We have to be narrow in the area we wish actuaries to be accountable and to what extent.

I do not have the answer to your question, Mr Smith, but I wanted to emphasise that you are making an important point. This paper is giving us the right now to raise the point, being conscious of the legacy we are going to leave behind for the next generation of actuaries.



Models are everywhere: you model price; you model retention; there are all the financial models. We have to be mindful of what the regulators will ask us to be the guardians of in the future, and if we wish to carry that responsibility.

**The Chairman:** So you are saying to be careful what you wish for.

**Mr Poracchia:** Absolutely.

**Mr Kelsey:** If you get rid of your model, you get rid of a risk. Therein you produce another risk. What is the risk of not having a model? You could be just taking that away. Some of us came into this industry when there were no models. It is a balance.

**Mr White:** I think to your question “What can we say about the model risk?” the honest answer is: “We do not know”. I think Mr Smith’s points are highly relevant, though, to discuss the things that possibly could be there.

When we talk about opportunities for actuaries, modelling and so on, it is critical to be aware of our lack of knowledge of the system in which we are advising. We feel moderately comfortable that we understand insurance, pensions, etc.

At a meeting here the other day, we were talking about the possibilities of being involved in valuing intangible assets of companies, which meant being an investment expert. I can think of little that would be more dangerous than that.

I agree with the speaker who said that we have to be careful. When we give advice in the areas being discussed here, we should recognise that we are not the only ones making assumptions. We have to be clear what they are; and our clients, the people we are working with, are often the right people to make those judgements, not us.

**The Chairman:** One of my objectives is managing model risk and uncertainty and presenting it to the board. Before I read today’s paper, I relied on two other excellent sources of ideas. One is called the “Comprehensive Actuarial Risk Evaluation”, which was a paper published by the International Association of Actuaries in May 2010. The other is a document by Standard and Poor’s entitled “Methodology for Assessing Insurance Economic Capital Models”, again published (I think by coincidence) in May 2010.

This paper provides a systematic framework for model risk management. It also addresses the thorny issue of expert judgement and cultural issues.

As regards the bottle of Champagne, I liked Mr Kelsey’s “obsolescence risk”; I liked Miss Strudwick’s “models get real”; I liked Dr Pryor’s “terrible trust in technology”. But I am going to go for Mr Kaye’s suggestion, not because of the phrase he came up with but because of the phrase he engendered. He said “model culture risk”, which I liked. Then in the following discussion we had “model robustness framework” and “model resilience management”. I quite liked those phrases because they are positive as opposed to defensive.