British Actuarial Journal, Vol. 16, part 3, pp. 701–721.
doi:10.1017/S1357321711000250

Measurement and modelling of dependencies in economic capital

Abstract of the London Discussion

[Institute of Actuaries, 10 May 2010]

Contact
Richard Shaw, F.I.A., E-mail: richard.shaw@hcr-uk.com

This abstract relates to the following paper:
Shaw, R.A., Smith, A.D. & Spivak, G.S. Measurement and modelling of dependencies in economic capital. *British Actuarial Journal*, doi: 10.1017/S1357321711000249

Mr G. Spivak, F.I.A.: (Introducing the paper) There have been many comments in the popular press about correlation matrices, in particular relating to the mortgage-backed securities market. A typical comment was from Nassim Taleb who said that everything that relied on correlations was charlatanism. I will open the meeting with a couple of words on diversification and correlation and why they are so important.

The most popular risk measure used in banking and insurance is the one-year 99.5% Value at Risk (VaR). For example, under the UK's Individual Capital Assessment (ICA) regime and Solvency II, an insurance company needs to hold enough capital such that there is a probability of 99.5% of survival over a one-year time horizon, or, in other words, the probability of insolvency over 12-months is no more than 0.5%.

Not all risks the company is facing will cause losses at the same time. Some areas of business may experience extremely high financial losses whilst others experience average losses, or even profits – an effect known as "diversification". Many firms calculate the capital requirement for each risk in isolation, ignoring the effect of other risks. The effect of diversification is visible when the overall capital required for an insurance company at the 99.5% level is less than the sum of the 99.5% individual capital amounts for each risk. The extent to which the aggregate 99.5% capital differs from a straight sum of the 99.5% individual capital amounts is a measure of the level of diversification between risks.

The diversification benefit depends on the level from which we started aggregating. If we started aggregating at the "equities, property, fixed interest" level we are going to get a different diversification benefit than if we start aggregating at the "UK equities, US equities, Asian equities, UK fixed interest, US fixed interest, Asian fixed interest level".

The modelling of dependencies and calculation of the overall diversification benefits goes beyond a pure bottom line insurance company impact. With the advent of the Internal Model Approval Process (IMAP) under Solvency II, and company desires to gain approval of their economic capital models for computing the SCR, various other criteria suddenly become important. CEIOPS's Advice for Level 2 Implementing Measures on Solvency II (formerly CP56: "Tests and Standards for Internal Mode Approval") states:

'Use Test (Section 3 of the CEIOPS Paper) Senior management shall be able to:

1) demonstrate understanding of the internal model and how this fits with their business model and risk-management framework.
2) demonstrate understanding of the limitations of the internal model and that they take account of these limitations in their decision-making.'

The timely calculation of results is essential. Management will need to ensure that the company avoids significant time lags between the calculation of model output and the actual use of the model in practice.

The extent of these requirements is not 100% clear. We do not think that the management will need to understand various types of copulas and how these are used within the company's internal model to calculate capital, but they probably need to understand whether the company's model takes tail dependency into account, how they model correlations under normal market circumstances compared to dependency between risks in stressed circumstances.

Another area of CEIOPS's advice is the Statistical Quality Standards. We have focused on the "Adequate system for measuring diversification effects" and in particular on paragraphs 5.245 and 5.246 of the paper. There should be meaningful support for claiming diversification effects that includes:

a) empirical/statistical analyses
b) expert judgement of causal relationships
c) or a combination of both.

In our paper we focused on how a company could go about meeting these requirements.

It does not propose one "best" way of modelling dependency. It discusses the advantages and disadvantages of each approach. Dependency is a very complex area of economic capital modelling allowing for a wide choice of model types and approaches to parameterisation. Issues that arise over a typical 12-month modelling time horizon are compounded using a multi-year model. A single correlation coefficient is often not enough to describe the dependency between risks in more extreme scenarios. In these scenarios the distribution-based copula approach to modelling dependency can be more meaningful.

Estimating a correlation coefficient can cause serious practical difficulties, including spurious relationships, availability of data and technical constraints. A simple scatter plot is likely to lead to a different interpretation than analysing the same information as an historical time series. A correlation matrix can cause a lot of technical issues, including positive semi-definiteness, high dimensionality and filling in the missing terms.

Using copulas can address some issues with correlations:

- Correlation is a single parameter, copulas introduce the distribution-based approach to dependencies;

https://doi.org/10.1017/S1357321711000250 Published online by Cambridge University Press

- Copulas allow a separation of the modelling of individual risks from the modelling of dependency between them;
- Copulas allow the direct modelling of tail dependence, although calibration based on past data suffers practical difficulties.

The benefits of greater flexibility need to be balanced with the difficulty of estimating a larger number of parameters. If copulas are used then the selection of an appropriate copula and its parameters should be based on sound analysis and judgement. However, there are considerable issues in trying to parameterise heavy-tailed copulas and so a pragmatic approach is often called for.

Copulas do not model the change of dependency structure over time, in particular, at different points in the economic cycle and, in the case of non-life insurance companies, the underwriting cycle.

A company needs to be extremely careful if it is using higher correlations within the variance-covariance framework as a substitute for tail dependence and copulas. The choice of correlations should not be based on the notion of a prudent margin in the absence of any analytical work underpinning the assumptions made.

Measuring correlations in key regions using half-space depth is a promising alternative. It allows the construction of:

- Combined stress tests, where more than one risk materialises at once. The intention may be to construct combined stress tests at a given level of confidence. This requires a generalisation of the percentile concept to multiple dimensions.
- Two-way correlation models, where different correlations apply between pairs of risk drivers according to the direction of a firm's exposure to those risks.

Both concepts do not naturally arise from the more traditional approaches to dependency modelling such as copulas.

Causal modelling, unlike other models covered earlier, allows modelling dependence between risks in a more direct and intuitive way. The paper shows how such methods can be used in non-life insurance or CDO models in the investment markets. The challenges using this approach include modelling the "stressed" states of organisations or capturing reinforcing effects. As in other types of dependency modelling, sparse data makes it difficult to evaluate objectively the merits of competing approaches.

Unknown contributor (opening the discussion): The idea of a diversification benefit or diversification credit is linked with the granularity of the calculation. It is meaningless if the granularity of the calculation is not stated. A diversification benefit could be misleading if it gets anchored in modellers' or managers' minds that 30% or 40% is in some way reasonable as that percentage on its own does not mean anything.

As the portfolio size increases, or the granularity of the calculation increases, you have an increase of diversification benefit. But at the same time you have a weakening in the correlations that are being observed to an extent that they may not even be detectable. As you increase portfolio size, the sensitivity of your aggregate capital to the correlation increases a lot. A situation can develop where your capital is, to an enormous amount, driven by your dependence assumptions because the calculation is highly granular yet you cannot specify the correlations because they are so weak.

This is one instance where an increase in the detail of your calculation does not increase the accuracy of your outcome, but rather the opposite happens: where your output is completely swamped by a model error.

Mr C. Keating (visitor, Finance Development Centre): Perhaps one area where the paper is sparse is that of causal dependency: there should have been something in there on Bayesian networks.

I wish to make a few comments about diversification and dependence. This is an area which appears to be misunderstood. A business can be diversified internally, like a bank; but banks are not actually a diverse group, they are systematically the same.

Diversification in the context of business lines within a bank, otherwise known as the economics of scope, is not the same as the idea of diversification across assets in, for example, a Markovitz mean variance portfolio. Failure of one line taints all of the others within a firm, which leads one to suspect some of the economic capital allocation models that are around, particularly in the presence of reputational risk.

Therefore diversification is a much broader topic than treated in the paper, and dependence is an even broader topic still. The very idea of specialisation of production in economics induces mutual dependence among an amazingly large network of people.

Diversification within an insurance business is the dependence, or the coupling, between the different sectors and parts of the production process. The looser the coupling, the more flexible the business is; the more rigid the coupling, the less flexible and the more prone to shock and catastrophic failure a business is.

Mr O. J. Lockwood, F.I.A.: Section 7.1.2 considers whether the risk factors used in an economic capital model should be the underlying economic risks, the business units or a combination of both. I support the first of these approaches. Using the business units as the risk factors is likely to result in dependency relationships between the risk factors that are unstable, because the risk exposures of the business units will change over time, whether through a conscious change in strategy or otherwise. Using a combination of both as risk factors is likely to result in an excessive number of parameters. The paper notes that using the economic risks as the categories creates difficulties in ensuring that the risks are defined consistently across the business, but this is an issue that can be kept under control through skilled management.

I also wanted to comment on the requirement for correlation matrices to be positive semi-definite (PSD). Section 7.7 suggests that, having derived a correlation matrix that is not PSD, one should apply one of a number of algorithms to adjust it to a matrix that is PSD. I would question whether this recognises the fundamental reason for the requirement for the matrix to be PSD. If a correlation matrix is on the boundary of being PSD, then it is saying that there is some linear combination of the risk factors that has zero variance. An example would be a model where the sum of the equity return, the property return and the change in interest rates is a constant. Such a model materially misrepresents the real world and there is a need for a method of estimating correlations which produces a PSD matrix in the first place, rather than for the correlation estimates forming a non-PSD matrix having to be adjusted to a PSD one. My solution is to derive a correlation matrix as the correlations of a set of time series. This will automatically produce a PSD matrix. Simulated data would be used where information about a risk factor is not available over the full historical

https://doi.org/10.1017/S1357321711000250 Published online by Cambridge University Press

period of the time series, or where the historical period does not include certain events which it is considered should be allowed for in the assessment of correlations. It would be possible to make the methodology underlying the derivation of this simulated data more transparent than the rationale behind a large number of subjectively assessed entries in a correlation matrix that is not PSD, together with an algorithm for converting the non-PSD matrix into a PSD one.

The issue of tail dependency potentially arises both in aggregating, for example, the equity risk distributions across business units to derive an aggregate equity risk distribution, and in aggregating the different economic risks. Copula-based approaches are likely to add significant value compared with the variance-covariance approach. There is strong empirical evidence that the dependencies between different equity markets tend to increase under adverse market conditions, and the quality and quantity of data available to calibrate a copula model are relatively high. In the case of aggregating the different economic risks, however, there is likely to be only a relatively small number of pairs of risks that both have adequate data to permit an accurate quantification of tail dependency. Equity risk combined with interest rate risk is the main exception to this. Rather than refining the approach to modelling the dependency between equity risk and interest rate risk, it is likely to be more valuable to refine the approach to modelling interest rate risk by considering interest rate movements other than a parallel shift in the yield curve. Insurance companies generally seek to hedge their exposure to interest rate risk and it is likely that the residual risk exposure will be positive at some terms and negative at others, and the type of yield curve movement that has the most adverse financial impact will then not be a parallel shift. Such more complex movements in the shape of the yield curve are less likely than a parallel shift to be strongly correlated with equity returns. The arguments I have outlined do not justify ignoring tail dependency in aggregating the different economic risks, but they do support the use of a relatively simple approach to allowing for it, such as the *t*-copula which has a single degrees-of-freedom parameter to control tail dependency.

Mr H. D. Sutherland, F.I.A.: The key issue that came out of the paper is that, if you are going to adopt the copula approach to modelling dependencies, then copula selection is non-trivial. There is a wealth of work to be done, and issues to be addressed, in deciding the appropriate copula to use.

The authors suggest that a Gaussian or t-copula approach might be the most appropriate. My understanding is that it is inherent in this approach to have zero tail dependency. The Gaussian copula seems to be inappropriate. The t-copula is symmetrical, so it treats upside and downside risk in exactly the same form. That seems to be inappropriate for the issues that we are looking at when we are constructing economic capital models. That leaves us with the Archimedian explicit copulas such as the Gumbel and the Clayton.

The paper does not help us decide which of those copula models is the appropriate one and how you can then go on to parameterise it so that you can use it in an economic capital model.

Mr D. Simmons, F.I.A.: My point relates to the use of the term “diversification benefit”. We should not be using this term as it has the potential to cause confusion when we are communicating within organisations, and to senior management in particular. As others have said, the measure that is typically referred to by the term “diversification benefit” is not really a benefit, it depends on the level of granularity you are starting from.

Rather than talking of “diversification benefit”, we should talk of “aggregation”. Aggregation is what we are actually doing, and this paper is about methods appropriate for aggregation. We should

https://doi.org/10.1017/S1357321711000250 Published online by Cambridge University Press

be communicating to our audiences and clients what we are aggregating (marginal distributions or different types of risks), what we are assuming when we are aggregating (correlation coefficients, causal dependencies, copulas, etc.) and, in particular, the impact of those assumptions.

If the focus is going to be on a 99.5% Value-at-Risk measure, then we should be able to explain the impact on this measure of making different aggregation assumptions. So discussing aggregation should be the line of communication, and we should stop using the words "diversification benefit".

Mr Keating: In the context of economic modelling and insurance companies and the regulator, one of the things to remember is that diversification, in the context of an aggregate systemic shock, does not help at all. For example, if you have an aggregate shock to liquidity, everything goes down, so the diversification within a business is immaterial.

One of the biggest weaknesses of correlation-type approaches is the absence of time in the model, whereas within causal dependent type models there is explicit treatment of time.

I cannot see how you can discuss the management of the business without considering the time part of the process.

Mr V. Knava, F.I.A.: Where I have seen the copula approach used in modelling, for example, in the case of one-year Value-at-Risk, it is a case of specifying a one-year marginal distribution along with a copula to aggregate all these distributions together. For insurance capital modelling, you might also want to know, not just the end of year investment result, but also the path along which you got there, in case you want to allow for things like management actions. Have any of you seen copula-type approaches being used in any sort of multi-period or multi-time setting?

Mr A. D. Smith (student): If you want to measure 12 months rather than one year, then you essentially have 12 times as many variables, and those different months will be dependent on one another. When modelling at more frequent intervals or modelling over time – this answers some of Mr Keating's questions as well – it does not change the nature of the problem, it just increases the number of dimensions you have got to calibrate.

Mr R. A. Shaw, F.I.A.: The causal modelling approach lends itself better to modelling dependency over time than the copula approach. However, causal modelling has its own challenges.

Mr Spivak: In our research of the literature a paper written by Mikosch, which criticises the use of copulas in financial modelling, states copulas are static. They do not describe time series modelling very well and you should use the proper time series rather than copula-type modelling.

Mr Shaw: On the copula approach, we have an example in the paper where we constructed a time series for two risks, A and B. We then went on to represent this information in two different ways. The first method was to construct a scatterplot with a view to then fitting a copula function to determine the dependency between the risks. The second method involved forming an opinion on the relationship between the risks from inspection of the time series data. The consequences of the copula approach is that one is effectively disregarding any time-dependency in the historical data. Even your marginal distributions are disregarding this time dependency. So if there is strong time-dependency in the data, then a causal modelling approach is likely to give you a different result than you would find with a copula.

https://doi.org/10.1017/S1357321711000250 Published online by Cambridge University Press

Mr R. J. Bowdrey (student): You mentioned parameter error and parameter risk a few times. Have the authors given any consideration to methods where you might deal with occasions where you do not have much data? For example, when you are looking at dependencies between catastrophe risks, for example? This is quite important when looking at capital.

Mr Shaw: With copulas the issue that you have is having historical data that is representative of the 'tail' dependency that one is trying to parameterise. In the ABC Insurance Company example, we considered a t-copula, with three different degrees of freedom, namely (i) ten, (ii) five, and (iii) two. This was not looking at historical data per se but very much seeing what the economic capital results are at different percentiles with use of these three t copulas. Furthermore we had economic capital comparatives of the Gaussian copula and the variance covariance matrix approaches. We were interested in how the results varied with changes in the assumptions.

One approach that I have tried in practice is to compare the economic capita results of a Gaussian copula with a t-copula of 5 degrees of freedom at different percentiles. If I used the same starting correlation matrix, the 99.5% or 1 in 200 year economic capital number would be higher through the use of a t-copula (five degrees of freedom) compared to a Gaussian copula. I reduced some of the correlation coefficients in the correlation matrix such that when a t-copula is applied the 99.5% economic capital number was equal to the economic capital number from use of the Gaussian copula but with higher correlation coefficients. Using the two different correlation matrices, I then compared the two copulas at the 90% (1 in 10 years) and 99.95% (1 in 2000 years). The t-copula gave higher economic capital at 99.95% but lower economic capital at 90%. I then formed an opinion whether these differentials were reasonable or not. If not then I would experiment with different degrees of freedom for the t-copula.

To summarise, parameterisation is very difficult. I would advise working around the boarders of a problem before firming up on assumptions based on the result sensitivities.

Mr Smith: People often draw a distinction between model error and parameter error. Model error is about having entirely the wrong set of formulae while parameter error is choosing incorrect parameters within the context of having the correct formulae. For these models what has struck me is how the model error dominates everything else.

So if, for example, you were fitting some copula by maximum likelihood, you could construct confidence intervals for the parameters. But you would be completely deluding yourself because these confidence intervals depend on the underlying assumption that your set of formulae is correct. The choice of the model makes a big difference.

Mr Shaw talked about varying the t-parameter and the t-copula. You can vary the number of degrees of freedom, but that does not give you a good picture of the range of results that could be got from any copula. It shows you the effect of changing one parameter in one particular family. The danger is you could be completely ignoring a whole load of behaviour that had not even crossed your mind. Mr Simmons made a similar suggestion, that one tests different values for the parameters that you put in. This is sensible but we also need to recognise that it is not sufficient. It is very likely you hard-coded your model to exclude all sorts of other things that had not crossed your mind.

In my opinion, the model error, thinking about the other possible models you could have constructed, is a much more challenging and potentially more alarming aspect than the error in calibrating these models.

Mr Spivak: My final comment is if you have not got any data, say on your catastrophe risk, you can ask yourself what is reasonable. If you are looking at the link between the catastrophe risk and inflation, you can ask yourself what are the chances of a catastrophe risk taking an extreme value, given that the inflation risk takes the extreme value. You can decide that from what you know about your business without looking at past data. That could give you your estimate for the tail dependence parameter, or for the copula parameter.

Mr Smith: This highlights the point that Mr Lockwood touched on, which is in real models you have a very small number of correlations which you can estimate using huge amounts of data, but the remainder of this huge correlation matrix is completed basically using guesswork.

If you fill a 30 by 30 matrix with guesswork, then it is very unlikely that it will be positive definite. As Mr Lockwood has said, it would be nice if you could create some fake data, whose correlation matrix you could then calculate and get something that was positive definite.

But going from mutually contradictory expert views, either to fake data or to a positive definite correlation matrix, are both pretty challenging enterprises.

For example, we were doing work with an expert that was totally persuaded that the correlation between interest rates and earthquakes was not the same as the correlation between earthquakes and interest rates. The reason, he argued, was he thought that you might have a change in interest rates resulting from an earthquake, but you are unlikely to have an interest rate change causing can earthquake. I agree with him that you are not likely to have an interest rate change causing an earthquake. But he was equally persuaded that these two correlations could be different numbers. However much I tried to say to him there is a mathematical fact that correlation matrices have to be symmetric, it was quite difficult to get any meeting in the middle.

We talk to experts in many different areas, but many of them probably are not experts in how to make a 30 by 30 matrix PSD. There is still a gap to fill. Most firms that I deal with at the moment are at the stage of having realised that there is a 30 by 30 matrix to complete and started to fill in some of the entries but are not terribly well prepared for the fact that when you get to the end some analyst is going to come along and say, "Er, this needs to be positive definite and it ain't."

Mrs K. A. Morgan, F.I.A.: To expand on what Mr Smith was saying, CEIOPS does quote in their advice to the European Commission "All models are wrong but some are useful". That emphasises where we are coming from. We do not think that models are the answer to the secrets of the universe, but they can be very useful to firms.

A key part of being useful is senior management's understanding of the model. People from the other regulators are here this evening ahead of a meeting tomorrow to draft some guidance for supervisors on how to assess senior management understanding.

One of the reasons we are so focused on this is because we had a look at what we could learn from bank models and how they were used. Our view is that the people who are using internal models have to understand them for their decision-making. This means understanding not only where the model works but also where it does not work, so the limitations need to be clear, and perhaps some of the alternatives.

The paper does suggest that using more simple methods would help communication. CEIOPS issued advice on proportionality, which people often misinterpret as meaning that small companies can do things more simply. That is not quite right. Proportionality means that what you do should be proportionate to the nature, scale and complexity of risks; proportionality is a two-way street so more complex risk will probably need more complex modelling.

Also, CEIOPS is keen to have models kept up-to-date. We have the foundation principle in our advice where we want the use of the model to create pressure for it to be improved. So we are keen that companies use techniques that are as up-to-date as is appropriate for them.

Section 4.1.4 shows aggregation techniques can give results that are either too optimistic or too pessimistic. That is why we are so focused on senior management understanding. This is important because these diversification effects can be huge, although this depends on the granularity of the model.

Another point is the paper states that in Solvency II firms can use the standard formula or an internal model. There is also the option to use a combination of the two and use a partial internal model for some of the risks.

One of things that we are thinking about in CEIOPS is the techniques we can use to integrate the results of the standard formula with the results of the partial model to come up with the final regulatory capital. We want the right model for the risk. While we might think that there is a right model that suits one company and other people copy it, we are looking for firms to use the right model in their own firm.

In section 6.7.2 the authors talk about calibration. But it read as though it was about validation as well. Validation is an important part of the Solvency II framework. Although all the tests and standards for internal models do combine into one framework, sometimes things are labelled as calibration when they are really validation and vice versa.

Finally, QIS 5 (Quantitative Impact Study 5): QIS 5 is going to be the final quantitative impact study before Solvency II goes live, and the FSA cannot say often enough that firms need to do it properly. This is not a quick job: it requires a proper multifunctional team to assess exactly what effect Solvency II is going to have on your company.

Mr A. P. Holtham, F.I.A.: Picking up on Mr Smith's comment about the 30 by 30 matrix and having to guess lots of the entries, as soon as you start looking at digging down into other product levels or looking at a group, you are not really talking about a 30 by 30 matrix at all but something like a 500 by 500 matrix. So whereas it is nice to think about a simple two-by-two example and how you might go about trying to derive that from data, in fact the task ahead for us is of a different magnitude.

My observation is that we are just not going to have enough data to let the data drive the assumptions and the copulas that we use. If you try deriving just a simple Pearson linear correlation from historic data sets, you can find the results are quite unstable, depending on what data sets you use, or they vary over time. So to get a consensus around, say, the correlation between equities and properties could be quite difficult.

As soon as you start thinking about dependency in extreme events, or tail dependency, then again the problems are a magnitude larger, because you just have not really got the data you need.

https://doi.org/10.1017/S1357321711000250 Published online by Cambridge University Press

There is going to be a very large amount of judgment involved in choosing copulas and setting assumptions. This is not necessarily a problem to begin with. You may choose a Gaussian copula or t-copula, because it is tractable and you can code it and make some slightly arbitrary assumption over the degrees of freedom. That seems to me a reasonable approach to take. But it is very difficult to foresee, in the near future, having the whole justification of your assumptions driven by data.

It might be a problem if we start trying to benchmark different firms against each other, which auditors and regulators will be trying to do at some stage. Under the ICA regime most firms have probably been using some kind of correlation matrix as their assumption set, whether it be in a variance-covariance framework or in some sort of Gaussian copula or maybe correlations embedded in some kind of economic scenario generator. So, at least slightly superficially, you have a set of 30 by 30 or 500 by 500 assumptions that you can use to compare one firm to another. My view is that a lot of these assumptions, where you have to allow for tail dependency, are set slightly arbitrarily. Then a process of benchmarking, industry benchmarking, audit review and FSA review, cause convergence into a more similar set of assumptions, although there is probably still quite a wide diversity.

If companies adopt diverse processes, whether it is using vine copulas or t-copulas or other types of copulas with different calibrations over these very large multidimensional numbers of risk factors, I struggle to see how it is going to be easy to benchmark sets of assumptions expressed in that way between one firm and another. That is going to be a real challenge.

Mr Shaw: Justifying assumptions is going to be a challenge in a Solvency II world for those companies that are seeking internal model approval for their capital calculations. When one reads "Tests and Standards for Internal Model Approval" (formerly CP56), one can see that justifying assumptions and validating expert judgment, etc., pose some interesting challenges, especially given that many correlation coefficients are set by so-called "expert" judgement in the absence of stable data sets.

It is very unlikely that the FSA would turn down every company applying for internal model approval and I would like to think that there is a degree of practicality here. If you took these rules to the nth degree, then no companies would get through.

To answer the question about the data, you are not going to be able to fully parameterise to a credible degree, in my opinion, a 30 by 30 correlation matrix. It is a question of materiality, focusing one's efforts on the more significant risk pairs with perhaps data or judgement for the rest.

Mr Smith: It seems to me that there is a lack of data for most of the risks that we are worried about. Where there appears to be plentiful data, there are questions over its relevance.

More worrying is that we go into this with a lack of self-awareness of what we are actually doing. You could be in the situation where you have got some assumptions that are made on the basis of data, albeit with an element of subjectivity about how it was done. When you are documenting what you have done, it is very easy to downplay the subjective element. It is easy to refrain from writing, "Here are the 30 other ways you could have gone about this and got completely different outcomes."

What people are looking for is something that looks like objective analysis. If you downplay all the other choices, it is not really objective analysis, it just looks like it. Similar issues actually arise from what we described as the herd mentality.

In a former existence I was working with some investment consultants. We would calibrate tail end models. We came up with risk premiums, that is the expected return on equities and property, and so on. We had a panel of fund managers. We would send round a survey once a year. The first year the results that came back were all over the place. For the next year they were a bit closer. By the third year, they were almost unanimous. What we realised was that they basically used the results of the last year's survey to calibrate the survey for the next year. But it gave the impression of improving consensus.

They could have easily seen this also with the ICA regime, where if you locked people in rooms with lots of equity data and asked them to come up with a 1 in 200 stress test then they would come up with a large number of different outcomes. Yet over the process of the ICA regime there was convergence. The convergence happened because people swapped notes and because people who had smaller assumptions than the rest of the herd had conversations with the regulator that suggested that they might want to be a bit more prudent.

The danger here is you see that is happening and say, "Gosh, there is obviously some progress in science, and where previously the people were all over the place, now they have come to a view and all agree on what the answer is."

If you do not have very much self-awareness, then you could easily think that this is a process of scientific research converging towards the answer, whereas all the science really tells you is that the incentives are such that people will converge towards an answer that is the same for everybody because it is not in anybody's interests really to be an outlier.

This is the difficult balance. You want to get the right answer for each firm. It is very difficult to judge that externally. What you can judge is whether one firm has done the same thing as another firm, and people are conscious of that. There is an issue here over the incentives. We would all like to see convergence. If you had real self-awareness, you might recognise that that convergence is the result of haggling, or note swapping, rather than the result of science. If you are realistic about how you got there, that is at least better than being under some illusion that science has spoken and proclaimed the answer.

Mr G. J. Mehta (student): The paper assumes that the model must fit past data. But analysing past data are you not making anchoring errors?

Actuaries have historically predominantly relied on past data to assist in predicting the future. Can we reduce the reliance somehow on the past data? This is not really addressed in the paper.

Mr Shaw: You have raised an important question, the issue of past data. Having specified the functional form of a dependency model, the next question to ask is what should the parameters of that model be? Then the natural approach is to say "We have historical data, let us try fitting the parameters to the data. Let us assume that the past is a good guide to the future". We do this day-in and day-out in all manner of actuarial analyses.

The question is: should we always do that? There is probably too much reliance on past data. After all historical data is just a snapshot from a range of possible historical outcomes.

Mr Smith: In terms of documenting what you have done, that is a whole lot more straightforward if you use past data.

https://doi.org/10.1017/S1357321711000250 Published online by Cambridge University Press

I have on occasions been in the position of being a maverick. In 2006, when corporate bond spreads were the narrowest they had ever been, if you had stood up and said, "I think we ought to be testing for a massive liquidity crisis where AA spreads went down 300 basis points", firstly, people would have thought you were crazy; and, secondly, you would never have got hired to calculate anybody's ICA for them.

We need to recognise that those incentives exist. You cannot make them go away. What you can do is interpret model outputs in the light of the incentives facing the people who put it together.

Dr P. D. England (Affiliate Institute): My first comment is about parameter uncertainty in relation to copulas and estimating parameters. Mr Smith mentioned using the Fisher information and doing something with that. That is a start, but probably a step forward is to use Bayesian Markov Chain Monte Carlo methods to get distributions of parameters. Obviously, if it is a single parameter problem, it is quite straightforward. If it is a multiple parameter problem, then you can use Gibbs sampling. This takes account of the dependency in the parameters, in addition to dependency in the actual underlying distributions. The computations get more complicated, but it can be done.

There was a paper on uncertainty with copula estimation by Jakub Borowicz and James Norman delivered to the International Congress of Actuaries in Paris.

The main comments I wanted to make were about defining what you mean by risk in relation to capital estimation.

In section 3.1.2 the paper says that there are three main components of an economic capital definition: risk measure, probability threshold or risk appetite and time horizon. I would say this is partially correct.

There are actually four because you have to apply a risk measure to something, and that is the bit that has been overlooked. You have to apply a risk measure to a risk profile. Under Article 101 of the Solvency II Directive, that risk profile is the distribution of the basic own funds over a one-year time horizon. The risk measure is Value-at-Risk, and the risk appetite is the 99.5th percentile.

One of the problems with the variance-covariance approach in the standard formula, which tries to apply correlation to capital numbers (instead of the underlying random variables) is that the risk profile itself is not well-defined.

From an internal model point of view, you would start at the overall capital number and, you would allocate backwards to risk types. But to do that you would have to define precisely what you mean by risk contributions from each risk type. I agree with a previous speaker that those risk contributions need to be defined very clearly.

It would be useful if the authors of the paper added the risk profile to their list, and comment on the suitability, or otherwise, of a variance-covariance approach to aggregating capital numbers.

Reference

Borowicz, J. & Norman, J. (2006). The effects of parameter uncertainty in dependency structures. International Congress of Actuaries, Paris.

Mr Smith: One of the most difficult aspects is the way the standard formula is laid out as essentially a series of algorithms. So you put a number in and a number comes out at the end. What has been missing is a clear articulation of the mathematical problems that these algorithms are purporting to solve. You can look at the algorithm and say, "I cannot make any sense of this." But one of the reasons that you cannot make any sense of it is because you have not been told where it comes from or what the problem is that it is trying to solve.

It is not quite as bad as you say. In our paper we have given examples of how one could back-solve for the parameters to put into a correlation approach and get roughly the right answer. I do not think it is completely crazy; but I do agree with you a firmer statement that says, "We are trying to solve the following statistical problem and that will result in this formula", would have been quite helpful. That is not a great deal of help to people who are actually calculating the formula. All they need to know is what the formula is. But for those of us who are fascinated by the modelling, I can sympathise with your frustration at the sometimes apparent opacity of what is in the standard formula.

Mr Shaw: The variance covariance approach has weaknesses, but it has strengths as well—its simplicity for one. In the Solvency II world you need to be able to communicate results to management and communication will be easier than will be the case with more complex dependency models.

What happens if I change the correlation coefficient from 20% to 30% for some risk pairs? We can see straightaway what the impact is. Furthermore, if you have a large worldwide insurance group, with 20 business units, within each of which there are 30 risks, so that you have 600 by 600 risk inter-relationships then a correlation matrix is likely to be a more transparent way of showing what these are.

The alternative of having a complex model with many dynamic interrelationships is more problematic to communicate. Communication of results and senior management's demonstration of the understanding of the internal model are important considerations in a Solvency II world going forward.

Mr Spivak: If I understand Dr England's comment correctly, I agree that the correlations should apply to the risk factors which drive the capital losses not to the capital losses. The same change in the same risk factors in equity levels may affect various companies to different degrees. If we are using a standard correlation coefficient, then it should apply to parameters directly observable from the market data; it should apply to the change in equity levels, not to the resulting change in the capital amounts, as in the current CEIOPS advice.

Dr P. D. England (Affiliate Institute): I agree with Mr Smith that it would be very helpful if we could have a mathematical description that actually made sense, so that when you add the capital numbers up you get something that does indeed correspond to the Value-at-Risk at the 99.5th percentile of the basic owned funds. The problem with the standard formula is that this has not been done, so there is no guarantee that it corresponds to what it is supposed to. Therefore it does not satisfy its own requirements.

Mr I. C. Marshall, F.I.A.: Like Mrs Morgan, I work in CEIOPS. We need to bear in mind that the standard formula is just a formula. It has some shortcomings.

https://doi.org/10.1017/S1357321711000250 Published online by Cambridge University Press

Firstly, the formula attempts to explain the risk profile of all companies across Europe, life and non-life; that same formula is obviously not going to be able to explain fully the risk profile of each firm. It is actually an estimation of the variance of the basic owned funds. The standard formula makes a number of shortcuts. In particular, it does not look at the change in risk margin which is one of the elements in the calculation of the basic own funds.

Secondly, it does not explicitly talk about the change in other liabilities. The standard formula is an approximation of the 99.5th percentile change in basic own funds.

Fortunately, the Solvency II framework also allows for the use of an internal model. With the internal model we have the same calibration point, in that the Solvency II regulatory capital must be calibrated to the basic owned funds Value-at-Risk over a one-year time horizon at the 99.5th percentile.

The good thing about the internal model approach is that it gives the firm the freedom to use the best approach that they can to get to that number. My point is that we need to bear in mind the standard formula is a "one size fits none" approximation to what the actual number should be, and the approach used in the standard formula does not need to be used in the internal model.

Mr J. A. Jenkins, F.I.A.: I should like to ask Mr Smith about his earthquakes and interest rates example. A more relevant example to me is pandemics and market crashes. A pandemic might cause the stock market to crash, but you would not expect a stock market crash to cause a 'flu pandemic. I have seen companies receive ICG on this point by not putting the correlation in.

You do not want to overdo it, so how do you allow for the fact that the dependency is not symmetric? You might want to allow for the fact that a pandemic could cause a stock market crash, but you would not want to be so prudent as to allow for it the other way round, when you know it is highly unlikely. If you are saying that correlation matrices have to be symmetrical, are you saying that you cannot use a correlation matrix properly to cover that situation, so you have to use a copula? Or can you cover that non-symmetrical situation adequately using a correlation matrix?

Mr Smith: I do not think we addressed that issue in the paper. I will try to answer it now.

There is a notion in statistics that correlation does not imply causation. One example is Mr Shaw's idea that if you go to sleep wearing your shoes, you are more likely to wake up with a headache. This effect being related to the common cause of having drunk a bit!

For many aggregation purposes the direction of causation does not fundamentally matter. What matters is the chance of these two events both happening in the same year.

To some extent, any statistical analysis, of necessity, is an analysis of whether things happen together or at different times, whether they tend to move in the same direction or not. But you can actually be agnostic about which causes which. I might have the bizarre belief that stock market crashes could cause pandemics, and although you might think me crazy for believing that, I might still come up with a perfectly usable capital model based on that belief. If we analyse the same data, we could end up with essentially the same model. In this case, if we consider causation, it is intuitively absurd to suggest anything other than the pandemic comes first, it actually does not matter for the purpose of much of this modelling what the causation is.

https://doi.org/10.1017/S1357321711000250 Published online by Cambridge University Press

Where it does matter is if you are building causal models, you are dealing with experts who do have views about the direction of causation, and so you need to be able to have a constructive dialogue. For calculating the variance of x plus y, you really need the correlation between x and y. You do not need to know whether x causes the change in y or vice versa. It simply does not matter.

Mr Jenkins: I agree up to a point, but if you do believe in the causation argument, are you justified in coming up with a lower correlation in the first place? In other words, putting the lower correlation into your model to allow for the fact that you genuinely believe that one will cause the other but not the other way round. Are you saying in your view you just cannot allow for that? It seems very odd to me to have to think about it in that way.

Mr Smith: You cannot start off with a correlation assumption of 50% and say, "Well, I am not putting half because the causation can only happen in one direction". That is not how we estimate correlations.

It is about how the experts can formulate their assessment of the situation. For example, it is not entirely unbelievable that a global economic situation could cause a pandemic, for example, if a government had an enormous budget deficit and could not afford the vaccines.

Mr Spivak: This issue is important if you are trying to assess a conditional probability of one risk dependent on the other risk, or something similar. But if all you are trying to do is to calculate the amount of capital required for the first event, and then for the second event, then you do not need to assess the probability that both these events happen together. As you are not concerned with joint or conditional probabilities, you should not worry about the causation effect at all. You are just estimating correlation as an average single statistic for two risks.

Mr V. Mirkin, F.I.A.: I should like to comment on the mathematical correctness of the standard formula approach. It is mathematically correct in a very specialised case when a company would be exposed to risks in only one particular sub-module, so ignoring the multistage aggregation.

Suppose all the risk drivers are elliptically distributed. For example, they come from a normal distribution and the capital requirements are a linear function of the risk drivers. Under those circumstances, the risk standard formula approach would give you exactly the right answer for the 99.5 percentile of the capital requirements.

Mr N. R. Bankhead, F.I.A.: Much of the debate seems to revolve around whether the primary objective should be accuracy or consistency. Actuarial work tends to rely on two separate approaches or philosophies. One is a view that the market knows best and the other is a view that history will repeat itself, in other words, what we have seen in the past is what we will see in the future. The latter is generally the philosophy that underlines models and we can model past history to various degrees of sophistication.

You can get as many models as there are actuaries, and most people could probably invent several models so you could get multiple numbers of models and most of them would give different answers. This may be inevitable but when the results constitute public information on which reliance is placed, this diversity can rightly be questioned.

A very real question is over whether history does repeat itself. Certainly, in the area of economics what we have seen over the years is that new economic theories have been invented to try to describe the circumstances of the time. Theory has evolved as history has evolved, largely because we have seen new circumstances. History does not repeat itself but produces new situations. This seriously undermines the reliance that can be placed on models and the degree of reliance that can be placed on past history.

I regard risk a bit like an iceberg. You can see some of it above the surface: the extent to which history does repeat itself. But there is a lot of it hidden under the water. I know with an iceberg you see about one-seventh of it floating; but the majority of it is hidden underneath.

By way of parallel, the same question should apply to risk models. When we look at risk, how much can we really see from what has happened before, and how much of it is unseen because circumstances will change and we will face genuinely new situations?

However accurately we try to assess risk, we will never manage to assess it with 100% accuracy. It is this very lack of accuracy that makes the question of consistency more important.

To what extent should we be looking for consistency as the main objective, perhaps with specific models for specific situations, or should it just be a free-for-all: anybody can use any model providing they can justify it?

Mr Spivak: Going back to the issue that the past data does not repeat itself. When you have a model and you do scenario testing or stress tests, whereby you test possible future scenarios, does the model allow you to forecast these future scenarios adequately given the assumptions you have made, even without any historical data?

Does the model make predictions in the right direction? Different companies employ different models just because they might expect completely different outcomes based on the same economic scenarios.

Mr A. Koursaris, F.F.A.: My comments relate mostly to causal models as an alternative to copulas or covariance matrix methods. I tend to use causal models because, in our work, it will help us to understand what the drivers of some of these risks are. Understanding these drivers will also help us to understand why they are dependent on one another, rather than just trying to find dependencies within some statistical data.

Some of the comments on causal models are fairly unfair and, in particular, the comments on overconfidence. The paper says if we develop very sophisticated causal models, it might lead us to be overconfident in using them. It is probably better to have a sophisticated causal model than a bad copula model that we have no confidence in.

The other points are on stress states and reliance within causal models. That is also unfair because a lot of these effects are also present when we use a copula but they are ignored when we combine these risks using a mathematical function, so we don't have to think about them. Whereas, making a causal model, we would actually have to think about what happened within our company in those states of the market.

There was a comment about it not being possible to go down to the lowest level of risk using causal models. But that is equally true for copulas and covariance matrices.

Regarding black boxes, it is true that we might create a causal model that becomes a big black box to the company. Using a t-copula to combine risk could also be quite a big black box.

On the point of the example on earthquakes or epidemics and how they relate to our economic risks, examples like that would be easier to calibrate using a causal model. For example, if this size of earthquake hits California, then I expect this much of a fall in the markets in America, but much less in the UK. In contrast, if I tried to combine my equity market risks after I have passed it through a market risk module with earthquake risk, I would find it very difficult to combine those sources of risk.

On the point of communication and understanding of models, I am not sure that, because these points are very hard to communicate to management, we should try to ignore them. I would rather see management trying to understand all of the main drivers of a causal model than a few of the copula assumptions that we use to combine them afterwards.

Mr Shaw: We are trying to remind people that when something is mathematically very sophisticated people quite readily believe the output. The causal models, like all other models, succeed or fail on the parameters that go into them.

For example, if we are looking at the relationship between natural catastrophe losses and reinsurance default, directionally we can say that there is a relationship. If we have an extremely large natural catastrophe scenario then we are likely to see an increase in the number of reinsurers defaulting, however it is a more difficult thing to parameterise such a relationship.

To reiterate, I know that copulas and various covariance matrices have their weaknesses but we just wanted to bring these points to the table so that they do not get overlooked.

Mr Smith: I agree with your points. Perhaps the best example is the modelling of defaults on collateralised debt obligations. These were the famous toxic assets, which were typically analysed by looking at the default risk of each underlying bond and then using a copula to model the joint distribution of the times to default of the bonds, and therefore work out the chance of lots of bonds defaulting at once.

A little-known fact about these models is that, if you are modelling the time to default for all these bonds using a Gaussian copula, then the chance of having a simultaneous default on two bonds is nil. The same happens if you shift to a t-copula or if you shift to any other of the standard copula formulas. There is zero chance of actually having two bonds defaulting at once.

Having two bonds defaulting at once is exactly what makes your collateral arrangements completely ineffective, because your collateral devalues at the same time as you have a loss on each of them.

So there is a really fundamental economic risk there which is not being picked up. There have been some fairly critical Wall Street Journal articles on this, one of which blames actuaries for having invented Gaussian copulas and hence bringing the global economy crashing down around us. The lesson people drew from this was, “You used the wrong copula!” whereas if you had used causal models you would hope somebody would have said, “Well, perhaps two of these things could default simultaneously because the same counterparty has failed, or some other reason”, and that could have been captured.

https://doi.org/10.1017/S1357321711000250 Published online by Cambridge University Press

So there are certainly some good examples in real life where causal models could have picked things up. The tragedy was, of course, that the models did not pick things up and we lost a huge amount of money. At least one could say that they could have done the job, whereas I am not sure that copulas could ever do that job particularly well.

Mr J. P. Ryan, F.I.A.: Mr Smith's comments are very helpful. It is the same issue that the data was not there for the CDOs because we had not previously had the same distribution systems for the banks, conglomerating all the risks together. Financial institutions, whenever they introduce a new product, generally aggregate risk in a way that we have not seen before. So, by definition, when anything new comes out it is not going to be in the data and therefore you cannot investigate it in that way.

The approach I like is to attempt to use a mathematical approach – copulas and whatever – to try and pick these things up. Then you do need to go back, using either some sophisticated modelling approach or just simply going through a "what-if?" situation. This might pick up some of these others things in a changing world.

The CDO is a very good, timely, example because modelling on historic data would not have affected the output because the important features were not in the data. It is important to think about whether the world has changed in such a way that these models need adjusting. This requires you to have a better understanding of the underlying business as well as the mathematics going through the model. An interesting argument for debate is: which is the more important, a thorough understanding of the business or an understanding of the model? The conclusion that I have come to is that they are both equally important.

Mr J. Instance, F.I.A.: The paper sets out the risks and the limitations of the models the authors have used. We have a new technical actuarial standards regime coming in. There is a TAS on modelling recently published and we have exposure drafts on insurance.

One of the key drivers of those standards is the explanation of the risks and uncertainties relating to the results that we are giving to the users of the information that actuaries provide.

There are risks and uncertainties in the results from any model. This paper is very useful in bringing out the risks in these sophisticated models used to determine the capital insurers require. In addition, the paper is extremely interesting and, for somebody who is not involved in this type of work very easy to read. We need to make it clearer to all users of the information that these are very complicated models and there is a lot of risk and uncertainty attached to the results that come out from them.

Mr A. N. Hitchcox, F.I.A. (closing the discussion): I attended the workshop before this meeting as well. I thought it was absolutely fascinating.

My job is that of a chief risk officer. I have to persuade capital providers and shareholders why I am doing what I am doing. In my economic capital model, the credit for diversification benefit is as big as my biggest individual risks, either from exposures or financial crises, so the loss of diversification credit is one of my biggest key risks.

This loss of credit could happen for one of two reasons. Either the risks do not diversify, because they happen at the same time, or the bigger worry is that I do not get credit for it in the supervisor's review or the findings of the rating agencies. Getting diversification credit right is worth hundreds of millions of pounds for most of us here. For those who work for large companies it is probably worth billions of pounds.

So I have to say that it is really helpful to have a paper that sets out the key issues so clearly and helpfully in just 120 pages. The authors have done us a great service and I recommend that everyone in the audience should read all 120 pages.

Sections 4 and 9 of the paper spent a useful amount of time clarifying the distinction between correlation and dependency. I am sure that many of us actuaries are going to be spending a lot of time in the near future explaining this difference to management.

Mr Simmons suggested that we do not focus on the term diversification benefit, that is, the reduction in risk. In my firm I also concentrate on the method of aggregation and how much the quantified risk increases as we add more risks to the portfolio. I have found that a very useful exhibit to show to Board members.

Section 5 of the paper sets out the different approaches to aggregation of the risk, and Sections 6 to 10 expand in detail on the mathematics and the data issues.

Mr Keating requested more information on the causal models, noting that this would help make it easier to capture the time-dependent effects. This is a point well-made.

Mr Lockwood supported using the underlying economic risks ahead of the business units as the key risk driver. My experience of working in a big multinational group is, after a while, head office asks you to do both. They ask you to look at the split by business units, then the split by risk categories, and then the two combined. What they are interested in is seeing if from different ways you get the same number or not. That is a really valid point from the point of view of head office.

On aggregation techniques, Dr England reminded us to be sure, if we are using the variance covariance approach, that we understand the underlying mathematical model, and to make sure that we keep clear the distinction between risk factors and capital amounts when we are using that approach.

On other technical issues, we discussed correlation matrices and PSD. All of us who are "dyed in the wool" practitioners know that we sometimes have to adjust big matrices that are nearly PSD to make them actually PSD. In the new world we are going to have to explain this to management: why we did it and what we did. That is one of the challenges I am not relishing!

Mr Smith also reminded us, on the topic of correlations, that there is not a two-way relationship between earthquakes and interest rates! Mr Jenkins questioned whether stock market crashes cause pandemics. I am sure those sorts of arguments will rage in boardrooms for many years to come. But the authors pointed out to us that often the capital required depended on the probability of the joint occurrence and was not affected by the direction of the causality.

On data issues, one speaker asked how we parameterise correlations and copulas for a risk with sparse data. Mr Holtham mentioned the 500 by 500 correlation matrix which needs to be parameterised in the tail, and the issue of benchmarking firms when they are all doing that. This was acknowledged as a very difficult area.

The authors pointed out that, in addition, the model error generally dominates everything else. They also cautioned on the risks of convergence of views with time driven by herd behaviour rather than genuine advances in science.

On communication of results, for me Sections 11 and 12 are the heart of the paper. If you believe that the Use Test is the acid test of justifying a capital model, then the best thing you can be doing is to be putting charts and tables, like the ones in the paper in Sections 11 and 12, in front of your senior management on a regular basis. When they see that a modelling choice can change an important number by £50 million or £100 million, then they will quickly get very interested in it.

Mr Koursaris mentioned the particular importance of the classification made in Section 11.1.7 on the different levels of granularity at which you provide the diversification information. I note his caution, but diligent readers of the “blue text” (in CEIOPS level 2 Doc 48 on standards for internal model approval) will hear echoes of Principle 1 with the Use Tests, that the Board must understand in which areas and on which entity, or hierarchy level, within the undertaking diversification effects arise. So getting that granularity issue right, or flagging it to management, is a very big part of your passing the Use Test.

Then we discussed how management might not need to understand the technical issues of copulas, if you are going down that route, but they do need to know that there are different choices, and that copula selection is non-trivial. We need to help management see the impact of the different choices. I am sure once we show them the impact of the different choices, then the questions will start coming as to why we chose one over the other.

Mr Sutherland asked for the paper to be extended to illustrate the workings of the Gumbel and Clayton copulas. If the authors have any more energy left, I would certainly like to see those results as well. So ideas for a further paper there perhaps.

Finally, on the topic of communication, I noted Mrs Morgan’s comment that the CEIOPS working party will be meeting and hope that they embrace the contents of Sections 11 and 12. Those sorts of charts look hard work; they are the sorts of charts that one day senior management will have to start getting their heads around.

Contributions on other topics included Mr Sutherland describing one of the text books as impenetrable, but this paper as transparent: I agree with him on that comment.

Mr Bankhead commented that risk modelling is like an iceberg. We can only see 10% of what is going on, and that the future will usually be different from the past because mankind is so inventive. But it is our job as actuaries to deal with that. We have to keep pushing forward.

Mr Instance reminded us of our duty to advise users of the uncertainties in our modelling.

Mr Ryan mentioned that a thorough understanding of the business is as important as understanding the mathematics, which I heartily endorse.

So, in overall conclusion, a really useful paper which generated a good discussion on a topic that goes to the heart of the economics of our industry. As we all know, insurance companies and other financial services companies are about accepting and transforming accumulations of risks into different sorts of portfolios, and one of our jobs as actuaries is to help management measure the value of these activities in the aggregate and at the margin. This paper is taking us a long way down that road.

Mr Smith (replying): There is not much to say in response because we have answered a number of the questions as we went along.

Professor P. J. Sweeting, F.I.A. (Chairman): It remains for me to express my own thanks and the thanks of all of us to the authors, the closer and all of those who participated in this evening's discussion.