

SESSIONAL MEETING DISCUSSION

Long-term stochastic risk models: the sixth generation of modern actuarial models?

Monday 7 June 2021

The Moderator (Mr K. Jennings, F.I.A.): Hello everyone and welcome to our session today, “Long-term stochastic risk models: the sixth generation of modern actuarial models?” My name is Keith Jennings and I am the chair of the Institute and Faculty of Actuaries (IFoA) Risk Management Board. I have the pleasure of chairing our session today. Our format today is a presentation followed by questions. Please submit your questions during the session using the chat functionality and we can cover them following the talk.

Our speaker today is Bill Curry. Bill (Curry) is a Senior Risk Management Actuary responsible for capital oversight and resilience testing at LV. He has over twenty years’ corporate and consulting experience helping firms to better understand their risk exposures through changing times. Bill (Curry) is passionate about the development of new modelling solutions and their application to give real business insights. And with that, I will pass over to Bill (Curry) for the session.

Mr W. R. Curry, F.I.A.: Thanks very much, Keith (Jennings), and welcome everyone to my presentation today. Initially, I would like to thank those who helped me with this presentation. I have had a lot of support from my colleagues at LV, from those at the IFoA and from the peer reviewers.

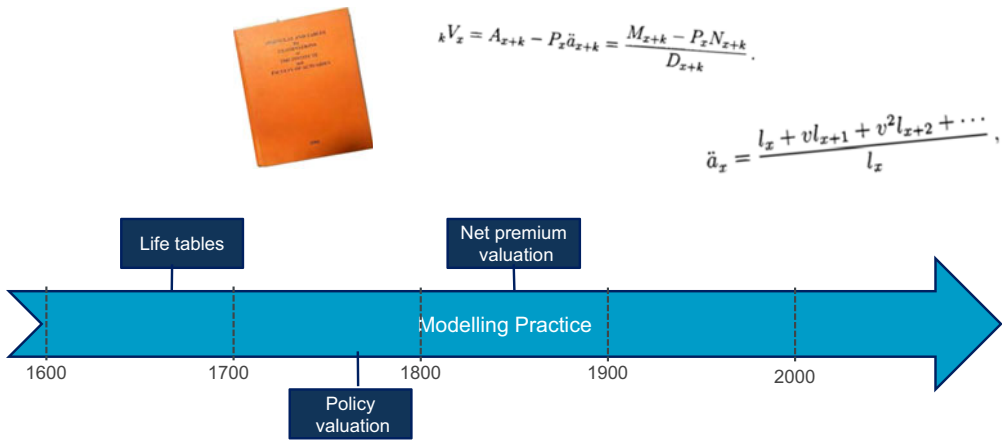
Without further ado, we will look at what is on the agenda today. We will talk about the history of some of the actuarial models that are being used, some of the market practice under Solvency II, some of the limitations of that practice. We will go on to look at how some stochastic long-term models can improve on those areas and look at some practical examples. The presentation is quite focused on UK life assurance, but, for those of you who might be working in other areas, I think a lot of the ideas and techniques are still quite relevant.

I will start off by talking about the history of actuarial models. First of all, I will discuss the kind of models that are based on commutation functions, formula tables and the use of manual calculations. These kinds of models have been around for several hundred years. Before this generation of models, there were perhaps the kind of models where people were just sharing costs for funerals, but it is really the timeline in Figure 1 that we are looking at here. This shows the real first generation of modern, scientific models.

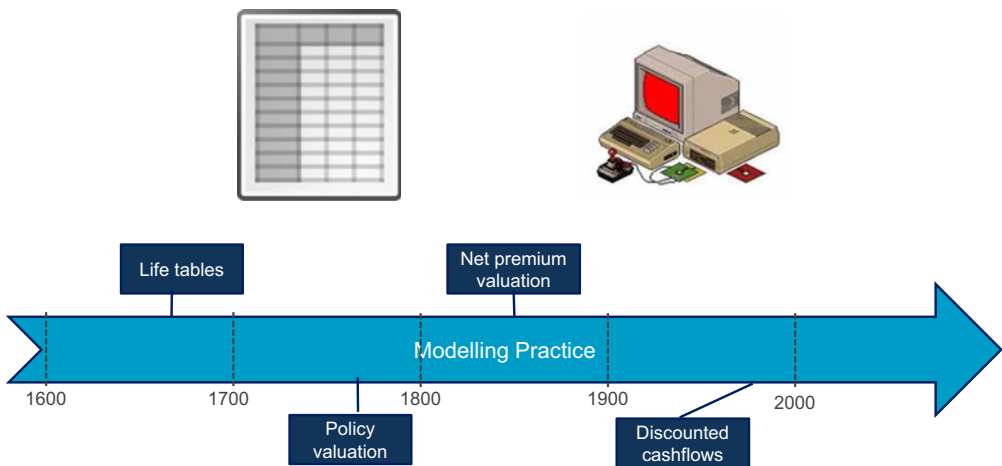
The timeline shows several important points. We have the development of the first life table in 1662, and the first policy valuation in 1762. The formation of the Equitable Life gave rise to a lot of modern actuarial theory. The first scientifically produced reserves also came from that time. So, this is the old approach that has been in place for several hundred years. I was working on models taking this kind of approach back in about 2002, and it represents the main approach that has been used since actuarial modelling started.

The next generation is about the advent of computer technology into life assurance. Computers started to find their way into life assurance perhaps in the 1980s. Computers are very good at producing lots of calculations, repeating them very quickly and without error. The advantage that

Generation 1 – Commutation Functions



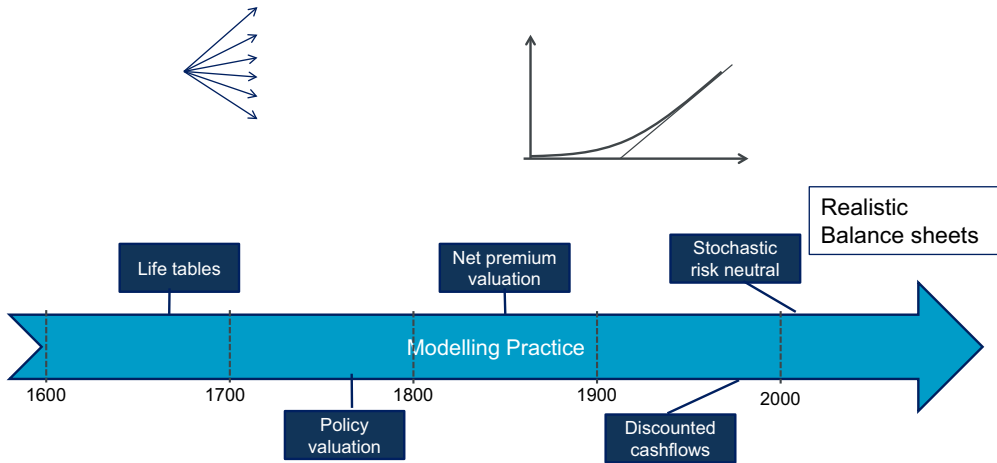
Generation 2 – Discounted Cashflows



this gives is to allow a more complex modelling approach, and this is where we have the advent of discounted cash flow techniques. The approach here, shown in Figure 2, is just to project all the future cash flows on a set of policies, discount them all back, to give your reserve, with normally a bit of prudence built into the assumptions to allow for any unforeseen circumstances. It is really the computer technology that has enabled this approach to come forward and be used. The benefit it gives over the previous approach is flexibility. Commutation-type approaches are really limited to quite a small number of simple products, whereas discounted cash flow models allow us to model much more sophisticated and complex products, and to understand in more detail how they work.

The next generation is market consistent models or stochastic risk-neutral models as illustrated in Figure 3. The key regulatory change here is the realistic balance sheet regime being introduced

Generation 3 – Stochastic Risk Neutral





in 2004, and with it, the market consistent valuation of guarantees. To deal with these requirements, you need models that can run simulations. The first generation of these ran about 1,000 simulations at a time to produce market consistent valuation guarantees, and would probably not have been feasible with the early computers we had in life assurance in the 1980s. They became possible when better technology later became available. When these models first came along, computers really struggled with them. A model I personally first ran in about 2003 had a four-day run time. That is how long it took to grind through 1,000 simulations. This sounds painfully slow by today’s standards, but that was where we were when these things first came in. I should explain, when I talk about these different generations of models, I am particularly focusing on the modelling practice and the regulatory framework. I am not so concerned about the invention of techniques. You often get a situation where a technique is invented and then only becomes the established practice at some later date.

Generation 4 – Risk Based Capital (Correlation Matrix)

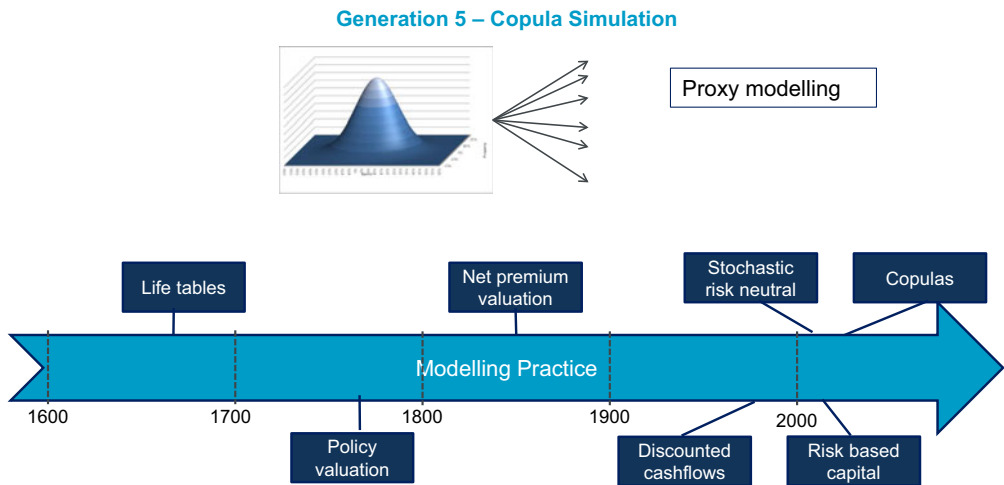
1	0	0.2	0.1
0	1	0.1	0
0.2	0.1	1	0
0.1	0	0	1

$$\sqrt{\sum P_{ij}C_iC_j}$$



Looking at the next generation, we are now onto the risk-based capital approach as illustrated in Figure 4. This was introduced at the same time as realistic balance sheets. The Individual Capital Assessments (ICAs) came in 2004. You could regard these models as part of the previous generation of models, but, I think, the two are quite distinct ideas. On the one hand, we have got a market consistent valuation of guarantees, and here we have got risk-based capital. They are quite separate. Here in the risk-based capital assessment, for the first time, we are looking at the capital requirements for all our risks individually and then bringing them together with an aggregation approach to get real-world capital requirements. This idea brings with it new techniques. As actuaries working on this, we became used to things like testing for positive semi-definiteness. We particularly became used to the aggregation formula shown in the middle of the top of the Figure. This formula is quite simple and easy to use. You just plug in the correlations and your capital requirements at the 1-in-200 level. It is easy to use, but it is subject to a number of limitations. Notably, it only works properly when you have linear loss functions, and then conditions have to hold around the probability distributions involved. We always refine techniques to get around limitations. Things like the single equivalent scenario technique would be an approach to try and get around the limitations of that aggregation formula.

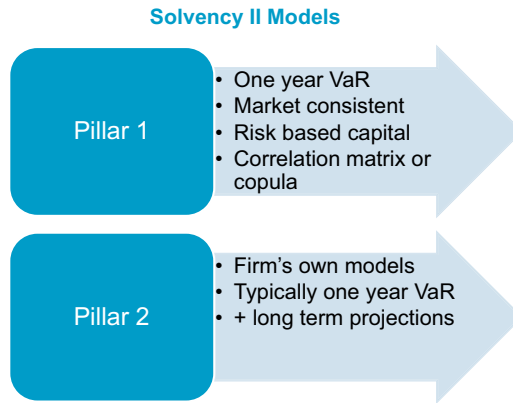
Figure 5 illustrates the next generation of models. This is the development of copula simulation modelling. Copula simulation modelling was invented in the 1950s. It came into the banking world in the early 2000s and into actuarial work in about 2008. The idea is to try to avoid the limitations around the aggregation formula. Rather than being limited to particular probability distributions and particular forms of loss functions, copula simulation modelling affords us the flexibility to model a variety of reasonable probability distributions and any form of loss functions. Therefore, it gives much better insight into risks. Just like with some of the previous generations, this requires another big upgrade in hardware capabilities, because we are probably running in the order of a million simulations. The PCs that were running risk balance sheets back in 2004 are not good enough to do this. There was also the advent of proxy modelling with this approach. As these copula simulation models came into use, there was also a requirement to have fast representation of changing balance sheets. That requires proxy modelling, which we as a profession have talked about many times in other presentations.



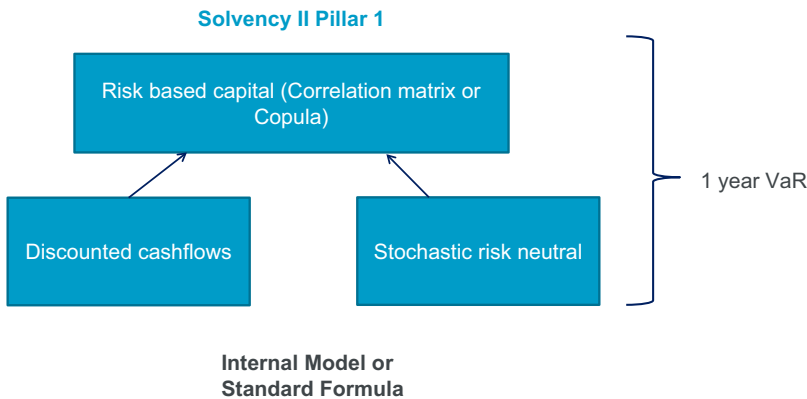
The key thing with all these different generations of models is that, very often, new hardware technology gives rise to the opportunity to have more specific models. Looking at some of the more advanced hardware and software packages on the market, we can run one million simulation

copula models now fairly easily. Models that were originally a big strain are now not too bad. We are perhaps in the position of being able to cope with a more sophisticated class of models again.

I will look at both the current regime and Solvency II. We have two main aspects to consider in Solvency II as shown in Figure 6. Pillar 1 is the regulatory peak, and Pillar 2 is about having your own view of the capital.

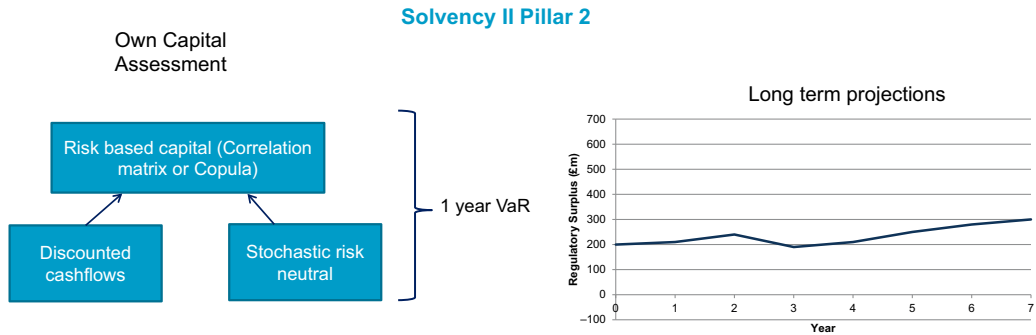


We look at those in more detail as detailed in Figures 7 and 8. Under Pillar 1, and this is the approach we are normally taking, you use discounted cash flow models to model things like your annuity business or other business. Then you use stochastic risk-neutral models for valuing guarantees. You bring those together in an aggregation approach. That aggregation approach might be a correlation matrix calculation if you use the Standard Formula. Or, if you are using an Internal Model, you are more likely to be using copula modelling, or you may be using correlation matrix modelling. We are using a lot of these previous generations of models. All these outputs are wrapped up in the one-year value-at-risk framework, which I will talk about more detail later.




Pillar 2, your own view of capital, is I think one of the strongest parts of Solvency II. There is a huge amount of flexibility here, and it allows actuaries to model capital in the most appropriate way for their businesses. For practical purposes, most firms have a Pillar 2 approach that is not too

dissimilar to their Pillar 1 approach. A firm’s Pillar 2 approach often uses the one-year value-at-risk approach as under Pillar 1, but perhaps with some different calibrations. There might also be different treatment of pension schemes or different treatment of Volatility Adjustments, etc. But a common approach to Pillar 2 is, essentially, based on the one-year value-at-risk framework.



Alongside that, there are also normally long-term projections. You might have a graph like that shown on the right-hand side of Figure 8, projecting regulatory surplus into the future in a deterministic manner, and looking at how that might vary depending on business plans. We often do this on a best estimate basis and then with various sensitivities or scenarios.

We now talk about the concept of the one-year value-at-risk that we use in Solvency II. The definition is shown on the right of Figure 9. It is really about two things: a probability and a timeframe. In a timeframe of one year, in this case, the value-at-risk is just the loss that corresponds to the given probability level. It is better illustrated by the diagram than by the definition. If your value-at-risk corresponds to the probability represented by the dark blue area under the curve in Figure 9, then the graph shows you what your value-at-risk level is. This concept was first published by J. P. Morgan in the mid-1990s, but some of the underlying ideas are actually quite a bit older.


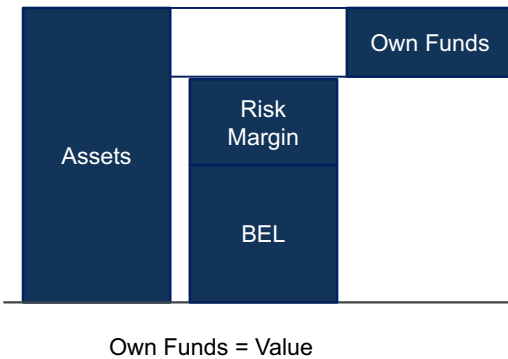


for a given time frame and probability, the VaR is the minimum expected loss on a portfolio for a given probability of outcome.

We will talk about how this fits under the Solvency II framework, and it is important first of all to talk about own funds. The own funds represent the transfer value of a business under Solvency

II. They are made up of assets less Best Estimate Liability (BEL) and risk margin as shown in Figure 10. The idea is that the BEL allows for the best estimate of the liabilities and also the transfer value of market-related liabilities or hedgeable risks. The risk margin represents the price of non-hedgeable risks, for example longevity, expense, Guaranteed Annuity Option (GAO) take-up risk and so on. Those are not included in the BEL. Therefore, if the business were to be transferred to a third party, the third party would expect to receive some kind of premium for taking on those risks. That is what the risk margin represents. The idea is that the assets less the BEL and risk margins are your own funds, and this is the value of the business. It is very important to understand that value. Risk margin is not a capital item. It is part of the value. Solvency II capital or the SCR is about protecting the own funds. It is the capital to protect the own funds with probability 0.005 over a one-year timeframe. That is essentially the idea of the Solvency II framework. The own funds are the value and the SCR is the capital to protect that value with a value-at-risk framework.

VaR within SII



SCR is the VaR capital to protect own funds for probability 0.005 over one year.

We will go on to talk about some of the limitations of the one-year value-at-risk framework. Before I do so, it is important to highlight that it has its strengths. It is very simple to understand and practical to apply. However, there are certain drawbacks as well, and I will talk through each of these in turn.

First of all, it tends to focus an awful lot on the 1-in-200 probability that we use under Solvency II. I know firms will certainly look at other probability levels as well, but capital, to a certain extent, will drive behaviour. Therefore, if the capital requirement is based on 1-in-200, then a lot of management focus will also be on 1-in-200. This could lead people to care a little bit less about events that are more extreme than 1-in-200. Perhaps people do not care enough about those events. However, those events can be important because, first, they could be the events that make the company go bust. Second, if the company does go bust, there is a big difference between going bust by a small amount and getting completely blown out of the water and leaving nothing left at all for the policyholders. Those two outcomes are very different, and the one-year value-at-risk does not really distinguish between them. Finally, if a policyholder has money in a life insurer for five or ten years, they might want management attention to be more focused on the 1-in-20 or 1-in-50 type events, and not the big, extreme, 1-in-200 events. The focus on 1-in-200 might take focus away from these less extreme events.

The next point to talk about here is pro-cyclicality. We use the one-year value-at-risk framework together with market consistency, and this combination can be quite prone to

pro-cyclicality. Take an obvious example. If the market starts crashing, an insurer will lose money on its assets. The increased volatility in the market will also spike up its costs guarantees, further worsening its balance sheet position. The insurer might also find itself in need of higher capital, and possibly a higher risk margin. All these things put pressure on firms to the extent that they might have to start de-risking and selling their assets, therefore sending the markets further down. In other words, this approach is quite prone to pro-cyclicality.

There are measures in Solvency II which try to get around this problem. First, the Volatility Adjustment is designed to protect firms against what is called artificial volatility. While it might do that, it is a move away from market consistency. It is ignoring what the markets are doing and using a different metric instead. Second, there is the symmetric adjustment mechanism used for equity risk. Under this mechanism, if the markets have been falling for a little while, then the equity stress gets reduced under the Solvency II Standard Formula. While that is useful for preventing pro-cyclicality, this mechanism has no actual economic basis. If the markets really have been falling for a little while, the chances are that there is high volatility, and you will probably need a bigger, not smaller, stress. The difficulty here is that we start with a pro-cyclical framework based on one-year value-at-risk with market consistency. In trying to prevent the pro-cyclicality, we move away from our principles of market consistency.

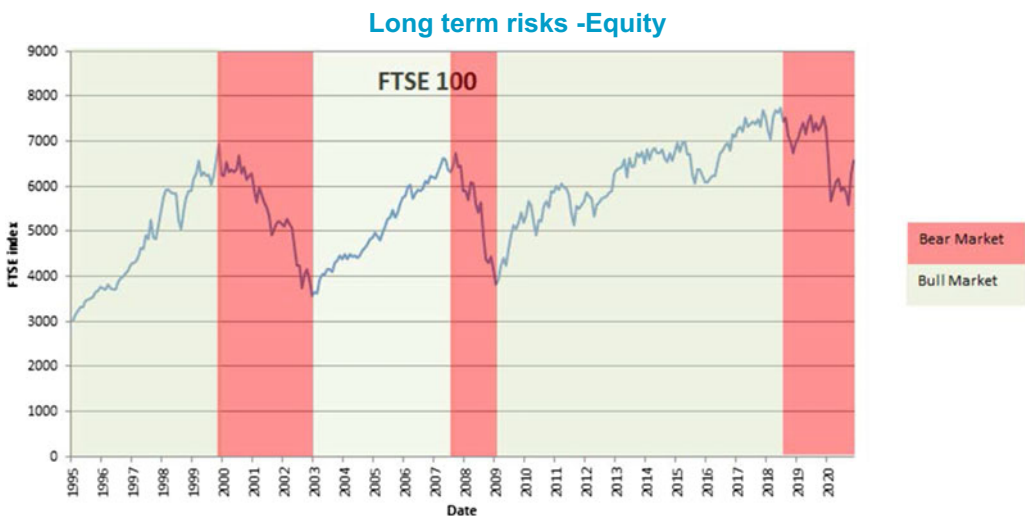
Next, the difficulty of calibration. Anyone who has worked on an internal model team will know exactly what I am talking about. It is generally very difficult to calibrate 1-in-200 risks, and also very difficult again to calibrate 1-in-200 level correlations and the way that these risks aggregate. One of our best-known risks in terms of data is probably equity risk. With the best will in the world, we have probably got about 50 years of decent, reliable equity data, if we assume that 1970s equity data is relevant to today's markets. With 50 years of decent quality data, do we really have a good view of what a 1-in-200 event is? Probably not. And it is worse for other risks. Some of the other risks like lapse risk and operational risk are generally very difficult to calibrate to a 1-in-200 level.

Next on the list is the point that diversification can increase capital. It is quite a shocking point. As an example, consider a situation with only reinsurance counterparty risk. Consider a business which reinsures all its longevity business away. Suppose there are two reinsurers on the market; they are identical and go bust with a frequency of 1 in 250. The direct insurer has a choice. It can put all its eggs in one basket and go for just one reinsurer or split the business 50:50 between the two reinsurers. Common sense tells us that it is better to split between the two reinsurers if all else is the same. The company would be diversifying its counterparty risk and should therefore be better off. However, they are not going to be better off from a capital perspective. Because these reinsurers go bust with a frequency of 1 in 250, if only one of them is a counterparty, then the direct insurer would have no capital requirements – happy days. But if the business was 50% reinsured with each of them, then the probability of a reinsurer default would be roughly 1 in 125, and the insurer would need to hold capital for it. This gives the counterintuitive situation where increasing diversification costs capital. It is a slightly contrived example, but the principle can hold more generally.

We are now going to look at market consistency challenges. The one-year value-at-risk framework that we use is all about having a base and stressed value of assets and liabilities. These values are supposed to be market consistent. However, this brings a lot of challenges where there is not a good market for the things being valued. A common example is property options. There is not much of a market for long-dated property options, so it is generally very difficult to get hold of a decent market value for them. There are also other examples like the Matching Adjustment. The Matching Adjustment is based on the idea of decomposing bond spreads into the liquidity premium and the default risk premium, and that is generally very difficult to do, because there is not a lot of market data available for this. The Matching Adjustment instead uses a retrospective approach based on historic defaults and downgrades. While that is probably the best approach with the available data, it is not especially market consistent. Another example is the risk margin.


The risk margin is meant to represent the transfer value of our non-hedgeable risks. There are lots of ideas on how to calculate the risk margin and how it could be improved. However, it requires finding a market price for something for which there is no market. That is generally very difficult.

On to the last point, the suitability of the Solvency II framework for long-term risks. The key point here is that life assurance is really long-term business. Policies might be on the books for decades. Do we really think a one-year framework is the right approach to model those? There are some examples in Figure 11, starting with equity risk. For equity risk, market traders talk about the tendency of the market to go on bull runs and bear runs. The bull runs are when the market is increasing, and the bear runs are where it is decreasing. The reason traders talk in this kind of language is that this is inherently how markets behave. They tend to have longer bull runs and shorter bear runs. It is about human behaviour. People like to ride their gains and cut their losses, so you tend to get long bear runs and slightly shorter bull runs. We can see examples here of some of the longer runs. Some of them go on for a lot longer than a year. We might ask, do we not care about the tendency of the market to go on these much longer runs when working under this one-year framework? A good example comes from the Dow Jones data back in the year of the Great Depression. The worst single year of the Dow Jones was 1931, when it lost 53%. That sounds quite, but this was in the middle of a three-year decline where it lost 90%. Do we not care about these longer-term runs where we could be losing a lot more? Are we only defending against one-year risk events?

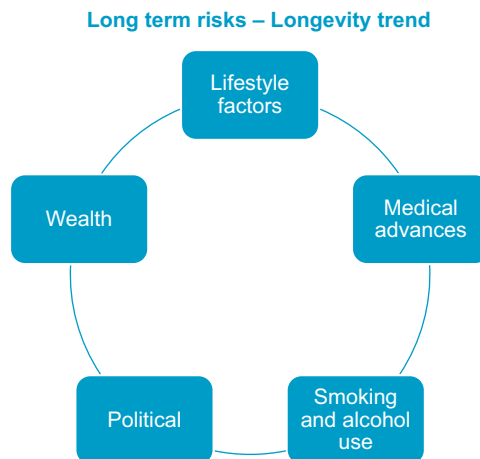


The Solvency II answer to this is that if you are monitoring your solvency at all times as you should, then if some of these bigger events come along, you should still, theoretically, have enough money to cover your own funds and therefore be able to transfer out your business to another provider. You may be able to de-risk it in the short term to prevent that in the first place, but if you have still got positive own funds, you should, in theory, be able to transfer your business to another provider and recapitalise. But this may be very difficult in practice in various situations. For example, if we are in the middle of a very deep decline in equity returns, you may not be able to achieve a transfer value that equates to this neat theoretical value of the own funds. And what if it is a systematic event? What if 10% or 20% of the industry is in trouble at the same time? Is there enough capacity in the industry to just sweep up the transfer values for all these firms? Another question is, what if one of the very largest firms goes down, is there really capacity for the rest of the industry to just sweep it up in this way?

We are going to look at another example now: interest rate risk. If you look on the far left of the graph in Figure 12, around 1985, a lot of firms such as the Equitable Life would have been selling guaranteed annuity business. This guaranteed annuity business is, of course, exposed to falling interest rates. We did not have risk-based capital at the time, but, if we did and if we had the current regime back then, a 1-in-200 interest rate events may have been, perhaps, a 3-4% fall. What that does not tell you is what could happen with very long-term declines. It does not model, or even attempt to model, what could happen with the long-term decline that has happened over decades. It is very important to think about these long-term outcomes because life assurance is a very long-term business. A lot of the business that was sold back in the mid-80s is still on firms' books now, and we are still exposed to some of these falling interest rate risks. Is this one-year value-at-risk framework really the most suitable framework for valuing these very long-tailed risks?



Another example is longevity trend risk. Many of you may have heard Stephen Richards' excellent presentation a few weeks ago about longevity risk under a one-year timeframe. One of the questions for the presentation was whether the one-year value-at-risk framework was really suitable for modelling longevity trend risk. My answer to that is that it probably is not. It is very challenging to model this kind of risk under a one-year timeframe, because longevity trend generally does not change very often or very quickly. The diagram in Figure 13 shows some of the



different things that can drive changes in longevity trend. These are things that happen very slowly over decades. The kind of example people use for a big change in longevity trend assumptive is a cure for cancer. This kind of thinking is a little outdated. No one really believes in a magic cure for cancer that is going to make cancer go away immediately. Medical progress in this area is more about individual treatments for individual cancers. But say we did have some magic pills discovered that were going to suddenly cure cancer. It is generally going to take several years for that to go through all the stages of development, through several stages of clinical trials, through the regulatory approval process, and then for it to be produced and widely used. It generally is not going to happen very quickly.

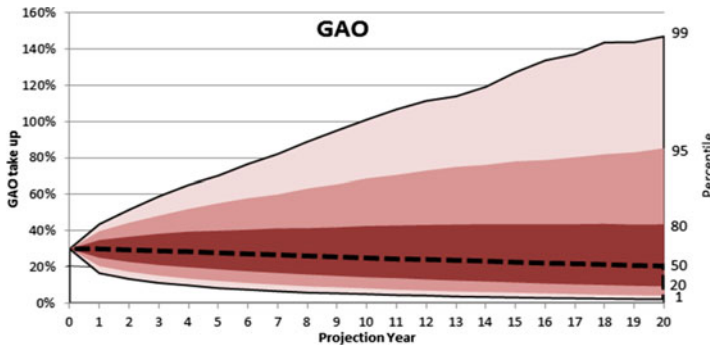
The counter-argument is, in the past year, has the world not responded extremely quickly to the threat of the coronavirus pandemic? While this is true, I think it is a little bit of a different argument. The coronavirus is a brand new, terrible threat that has come on the horizon and the scientific community has responded incredibly quickly to that threat. It is a different thing altogether to make quick progress against conditions that have been in place for years. As scientists and doctors have been battling against heart disease and cancer for years, it is difficult to see how a massive, sudden step change could happen in quite the same way. Also, if the discovery is of a preventative treatment rather than a cure, then the take-up rates might take a long time to build up. If someone came up with a new, cancer prevention drug then, I imagine, the take-up rates would start off not being particularly high because people are quite risk averse and tend to think, "Well, I am no worse off if I don't take it." Take-up rates might gradually build up over time with increasing public confidence. So longevity trend is another example of a risk that is very difficult to model under a one-year-time horizon. When I was putting this presentation together, I thought longevity trend was probably the ultimate long-term risk, but then, thinking about it more recently, climate change probably gives it a run for its money. Climate change is possibly the ultimate long-term risk.

I am now going to talk about the limitations of copula simulation models. First, I will explain how these work for those who are not familiar with them. We feed the copula simulation model with assumptions on the probability distributions of risks. We simulate different risk movements from these distributions. We then estimate losses from each of those risk movements. We might run, for example, a million simulations, and therefore have a million losses. From there, we can effectively build a distribution of losses and derive the SCR by taking percentiles from that distribution.

There are two main points to consider with regards to these models. First, they are single-period models. They model the start of the period, and what has happened at the end of the period – one year in this context. Second, they are based on statistical relationships between risks. By that, I mean that the assumptions put into the copula simulation model are things such as the correlations and probability distributions for the different risks. As everyone knows, correlation does not imply causation. What we are really modelling here with the copula simulation models is the joint likelihood of different events occurring. We are not modelling what causes what. We are just looking at the combined likelihood of different events happening at the same time.

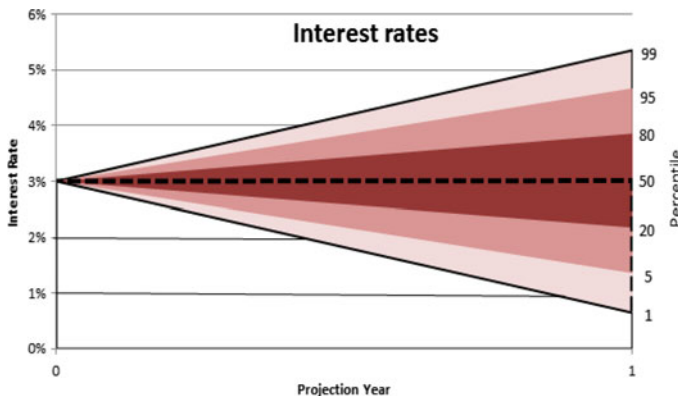
There is a problem with copula simulation models. In general, these models are very good at implementing the one-year value-at-risk framework, but the difficulty they bring is that we are very heavily invested in these models. They are very inflexible if we want to use a different time horizon. If I have my copula simulation model averaged over a one-year timeframe, and I need to work out what is going to happen in 20 years, what can I do? You might think I can just repeat my copula simulation model, but that is generally a bad idea. You can get something like the situation here shown in Figure 14. I will illustrate it with GAO take-up risk. If we just repeatedly simulate from the same distribution, the results get implausible very quickly. In this case, the GAO take-up rates exceed 100%, which is obviously impossible and makes no sense. A similar example on interest rates is illustrated in Figure 15. On the left is my one-year simulation. That might work quite well, and I might have calibrated it by fitting a probability distribution to historic data. If I repeat-simulate that over a longer period, the rates just keep on spreading out further and further. You get some very negative rates, and it all gets very implausible again. You cannot just repeat-simulate from these things.

Copula Models – extension to different time periods
GAO take up example



Starting take up rate = 30%
 Change ~ Normal(0, 19.4%²).

Copula Models – extension to different time periods
Interest rate example



You might think, if we cannot do that, can we instead just calibrate our copula simulation model to a different timeframe in the first place? If you wanted to know, say, movements in GAO take-up rate or interest rates over a twenty-year period, can we not just calibrate a copula model to that point in the first place? In theory, you can, but it is practically very difficult to do so. Initially, if you are looking at what happens to things like interest rates over a twenty-year period, that is generally very difficult to do by fitting probability distributions to historic data. You would probably only get a sensible result by looking at different sorts of models, models like time series. Copula models would generally be very hard to calibrate, particularly the correlations. If you are looking at how two risks correlate with each one based on twenty-year timeframes, the only way to reasonably attempt to get close to that is through a heavy amount of expert judgment and through heavy reliance on overlapping data, which has its own challenges. There is also the important point that it is not just where you get to that is important, it is how you get there. The path that takes you to the end point is very important. Consider equity risk, for example, over a twenty-year period. If I knew it had a fall of 10% over twenty years, that might be the final outcome. The other thing to look at is how you get there and whether there are big falls at particular periods or periods where the equity market is flat for a long time. How you get to the final outcome is very important,

particularly if you are looking at modelling management actions. We will talk about that a bit later. That is why these copula models are not very easily extendable. First, they are very difficult to calibrate to a different time period, and second, you do not solve the issue of path dependency.

For longer-term modelling, it would be better to use Cox–Ingersoll–Ross (CIR) models as shown in Figure 16. We can use these to very easily model interest rates over a longer period. The key things to look out for with this type of model is that we are really attempting to understand the way interest rates work rather than just fitting distributions to data. In the model here, we have a mean reversion term built in, and an assumption on the way volatility works – proportional to the square root of the rate. We are making assumptions on the way interest rates work and understanding interest rates, not just fitting distributions. At the end of the day, copula models are single-step models and use statistical relationships between risks rather than attempting to understand the true causal relationships.

Copula Models – extension to different time periods Interest rate example

Single Period Statistical model
Change in interest rate \sim Normal(a, b^2)

**Stochastic Differential equation model
(The CIR model)**
 $dr_t = a(b - r_t)dt + \alpha r_t^{0.5}dW_t$
where r_t is the short rate at time t
 a, b, α are parameters
 W_t is a Wiener process

We will go on now to look at why we might need long-term stochastic models. There are a number of different ideas as are shown in Figure 17. You can probably think of some others as well. We will start with risk management which is the most obvious objective. We talked earlier

Why do we need long term stochastic models?



about the limitations of the one-year value-at-risk for looking at long-term risks. Long-term stochastic models can get around this. Whether you are looking at interest rates or longevity, with a lot of long-term risks, you can avoid problems with the one-year value-at-risk framework and understand the underlying risks properly by using long-term stochastic models. You might think that, for company business projections, you will have a base projection and some scenarios already. Would these not give you these insights anyway? Stochastic modelling gives you a lot more than scenarios do. For example, stochastic modelling will tell you how likely events are as well as just what the cost of them are. It can also identify events that you might not have thought of before. With scenarios, you might pick two or three and test those. That does not give you any insights into scenarios that you have not thought of. Stochastic modelling can do that.

The next point in the Figure is strategy. Strategy is all about a balance of risk versus reward. You cannot really understand the risk element very well without some kind of stochastic modelling. We have got a practical example at the end of this presentation where we discuss the strategic side of things.

The next point is pricing. There are some applications in pricing. What is going to happen to the business over the longer term?


Estate distribution is an interesting issue. It is quite a familiar problem for many actuaries of having a with-profits estate and a closed fund that needs to be distributed to the members. The question is how to distribute the estate to these members in an equitable way. A market consistent framework does not necessarily work very well to answer that question. The reason is that if you have no plans to sell the business or transfer it out, and the value is all just to be distributed to the members, then the market consistent value does not really matter much. Long-term models allow you to look at groups of policyholders and the expected return that various members, or with-profit policyholders, would get, as well as the risks associated with these policies. This can help you come up with a strategy that is intended to be fair to all.

Last on the list is capital, with a question mark. Is a one-year value-at-risk framework the right capital approach to use, or is something long-term more appropriate?

We will go on next to talk about long-term stochastic models and what we mean by them. First of all, a long-term stochastic model is a model that looks at all your risks which are outlined in Figure 18. I do not mean things like strategic risks or the risk that regulations change, but rather the main risks that you might see in an internal model. A lot of firms probably do a lot of this kind of modelling already on the market risk side and look at how market risks might affect their future projections. However, that is not going far enough for me. You would not say your internal model is very good if it only looks at market risks, and nor should you think your long-term projections are very good if they are only looking at market risks. You really need to understand all the main risks for the business. You need to be able to, essentially, model each of these risks using


What do we mean?

Model all the key risks



What do we mean?

Use multi-step (time series) models to derive simulation paths for each risk




simulation time series as shown in Figure 19. For each of the different risks, you need to be able to simulate a time series of the appropriate values. Additionally, as much as possible, you need to use independent Monte Carlo simulations for different risks, and then build in the interactions between them through causal rather than statistical relationships.

Now, to illustrate what I mean by that, I will show you what I do not mean which is illustrated in Figure 20. First of all, this is how a copula model works. It is based on the statistical relationship between risks. The idea is you start off with independent random numbers. The copula provides the statistical interaction between these and gives you dependent, correlated random numbers. From there, you can get your equity and lapse distributions. That is the idea behind copula modelling. As I said before, it is all about statistical modelling.

What do we mean?

Relationships between risks based (as much as possible) on

- Independent Monte Carlo simulation
- Causal relationships



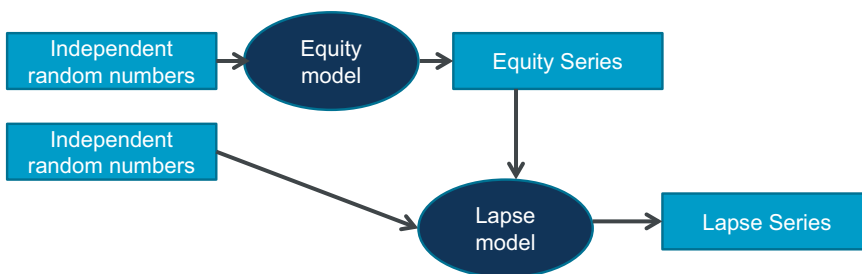
We do not want to do that. We want to do something more like what is outlined in Figure 21. This is an example of causal modelling. We have just two risks here, equity risk and lapse risk. We want to model the way that, if there is a fall in the equity markets, there may be an increased number of lapses as a result. This is a causal relationship rather than a statistical relationship.

Causal relationships can be much more useful than just looking at statistical relationships because they tell you what drives what, and can give you greater insights. In this example, you independently simulate an equity series, and the equity series feeds into the lapse model. The lapse model is therefore depending on some independent random drivers and the equity series and gives you your lapse series.

What do we mean?


Relationships between risks based (as much as possible) on

- **Independent Monte Carlo simulation**
- **Causal relationships**



A good example of a causal model is the Wilkie model from 1984. Figure 22 shows the structure of the Wilkie model, which tries to model some different market risks. All these different metrics are driven by independent Monte Carlo simulations, but with the interactions between the risks built in through causal relationships. A lot of these ideas have been around for a long time.

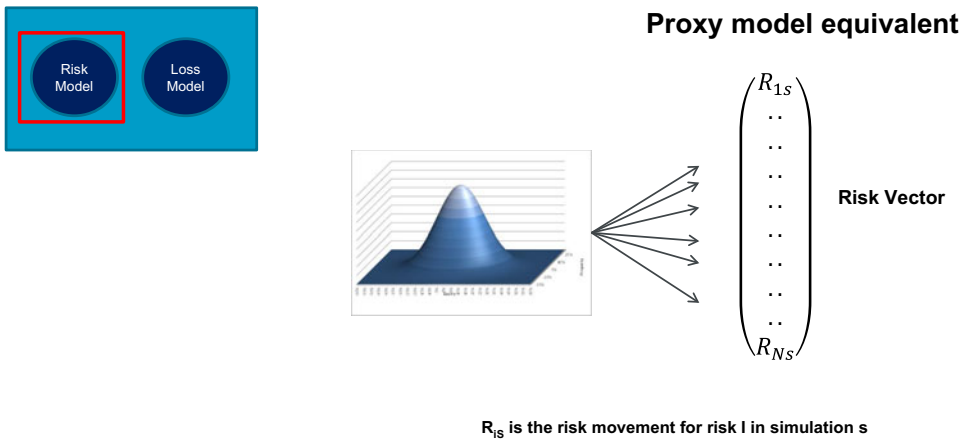
The Wilkie Model



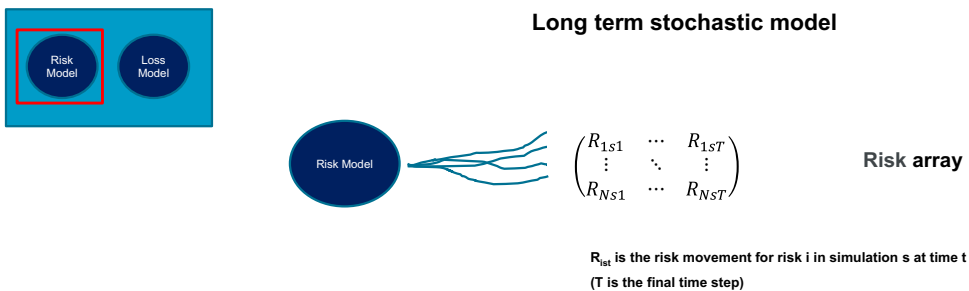
We will talk about the practical implementation, and there are two sides to this. There is a risk model and a loss model. The risk model is effectively the model of the likelihood of different events occurring, and the loss model is what these events do to your business, or what these events cost if they do occur. If we can think about these two things in turn, then we can understand the nature of the risks.

We start with the risk model. Figure 23 shows what a risk model looks like for proxy models used with copula modelling. I use this as an example because I know quite a few people are quite familiar with copula modelling. In copula modelling, you take a probability distribution and simulate a risk factor. The rows in the risk factor represent the different risks. You might be simulating a million different versions of this risk factor according to your probability distributions. In long-term stochastic models, you would add the extra dimension of time as shown in Figure 24. Here, in this array, the rows still represent risks, but the columns represent the time steps. You might have, perhaps, monthly time steps, and you would need to be able to simulate this big array of risk movements that represent all the different risks over all the different time steps.

Model Structure – Risk Model

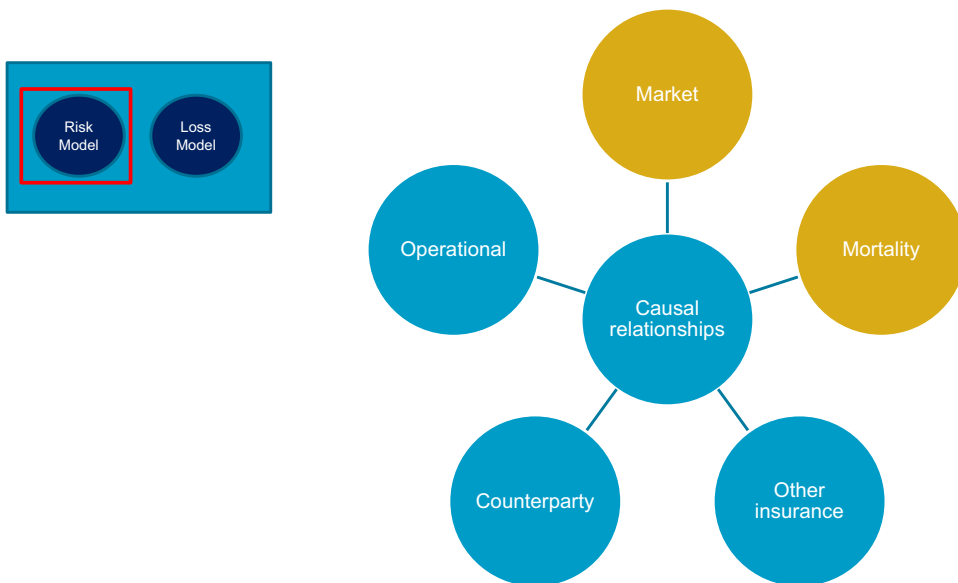


Model Structure – Risk Model



We can look at how this works for all the different risks, and at the heart of this is causal relationships as shown in Figure 25. I think it is important to use causal relationships as much as possible to understand the way that the different risks fit together. I will start by talking about market risk and mortality risk.

Model Structure – Risk Model



For market risk, real-world Economic Scenario Generators (ESGs) have been around for a long time. They are very good and very reliable. You can just use them and do not need to do anything new on that front. Similarly, with mortality (and I think mortality and longevity are the same risk in different directions), there are some great models that are very widely available for doing stochastic mortality modelling. A lot of firms already have some good stochastic longevity and mortality models.


For operational risk, I think most internal model firms will have operational risk modelled using scenario workshops, looking at various operational risk events and what they cost. If we need to turn that into a long-term model, it is really a case of thinking about how long events last. Is an event likely to happen over a particular timeframe? If an event happens once, does that mean it cannot happen again? The existing operational risk models can be adapted to work on a long timeframe with a bit of adjustment and a bit of work.

A simple way of approaching counterparty risks would be just to assume rates of default for your counterparties, assume recovery rates, and model defaults that way. A more sophisticated way would be to look at, perhaps, the three biggest counterparties and consider their own risk exposures. Are they exposed to, perhaps, US equity markets? Are they exposed to mortality? What are they really exposed to? Try and build in a causal relationship along the lines of, “If the US equity market has gone down significantly, then there’s a greater counterparty default risk.” I think that gives a much stronger approach to modelling than just using correlations to assume a statistical relationship.

The last category is the other insurance risk. This is going to include things like lapse risk, expense risk, GAO take-up risk, morbidity risk, and so on. Models for these are not widely available for firms at the moment. There is perhaps some new thinking and new modelling required here. I am going to give an example a bit later in this presentation about how we might do this, using lapse risk as an example, but we can do that for a number of other risks as well. I shall come back to that point.

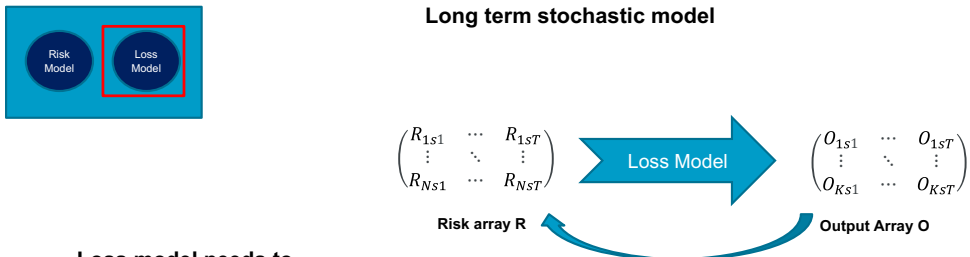
I will now move onto the loss model. Again, we are looking at proxy functions as shown in Figure 26. This is the way proxy modelling works. The loss function is a function that will take the risk factor and turn it into a loss. Figure 27 shows how this works for long-term stochastic models. You need to take a risk array, as was discussed earlier, and turn that into an array of outputs. The most obvious output here is the loss risk for time-dependent losses. These are things like economic variances, demographic variances, and so on. But in addition, you can also look at a number of wider business metrics, such as solvency and dividends. We will talk about this shortly. Note that outputs might in turn affect risks. For example, if one of your outputs is solvency and the solvency ratio is low, then perhaps more people will start lapsing because they do not feel comfortable with your firm anymore. Outputs might affect the risks.

Model Structure – Loss Model



Loss function gives the cost of any value of the risk vector

Model Structure – Loss Model



Loss model needs to use the array R of risk movements over time to produce an array of outputs over time

We now look at the output array in Figure 28. The first one I have put in with the biggest letters is the loss, because that is the most important one. We will come back and look at that shortly. If

Model Structure – Loss Model



you can model your losses, then you can get a lot of insights into the different risks. However, if you can model some of the other outputs, not necessarily all of them, then you can really understand your business in a very strong way. You can use these for strategy testing and for long-term projections of your business plan.

Liquidity is something I have not really touched on so far. Liquidity modelling is also possible with this kind of framework, and it is perhaps quite difficult to do with one-year value-at-risk.

We will look at the losses element that we talked about, and the reason it is so important. Losses are relatively straightforward to model. If you can model losses over different time steps, as shown in Figure 29, then you can calculate a run-off-based capital metric. We have two different capital metrics here. We have the existing value-at-risk metric on the left, and the run-off-based capital on the right. Run-off capital is effectively the amount of capital you need to pay off policyholders if you get an ebb of confidence. It is a very simple and intuitive concept, and it is, perhaps, a bit more intuitive than the way the value-at-risk works. I actually think it is much better from a policyholder’s perspective. If I am a policyholder, I am putting my money in a firm. I might be expecting to get some money back in 30 years’ time. I am investing in that firm. I want security that they are

Model Structure – Loss Model



VaR based capital ensures assets > BEL + RM after a 1inX stress

Run off capital is the amount to pay policyholders at Y% confidence

going to be around to pay off my policy. That is what run-off capital gives you. Value-at-risk capital does not give you that. It protects you against the big shock. It is saying that if a big shock event comes along, then the insurer might need to transfer your business to another provider for a hypothetical transfer value, including the risk margin. The run-off capital is a much stronger, much more intuitive way of thinking about things. It is something that might be useful for firms to look at in their Own Risk and Solvency Assessment (ORSA). They can look at the amount of capital they have got on a run-off basis, and see if that is a suitable amount to pay off all their liabilities.

The last part of this section is about practical implementation. These are some key ideas to think about with the design of long-term stochastic models. First of all, I always think that centralised design is really important. If you are buying a new mobile phone or a new car, the design of that phone or car will have a central component to it. There will be a central design team that has the authority to design the main parts of it. You cannot just have a free-for-all with everyone designing their own parts. You have to have a centralised point of design that owns the main design, which then delegates things out as much as possible. In insurance, you might have market risk experts dealing with the market risk. You might have longevity experts modelling the longevity part. You might have experts in the way the balance sheet works and the capital requirements, dealing with the losses side of it. The idea is you have a centralised design, but delegate things out as much as possible to the relevant experts to bring it all together in a single framework.

I said I would come back to a practical example. The following example is about lapse risk, but it could equally well be applied to some of the other risks, maybe expense risk or GAO risk. The presentation today is not meant to be a lapse risk presentation. It is just an aside to show what can be done in this area, and for that reason, it is a bit light on the calibration. It just touches on the theoretical elements.

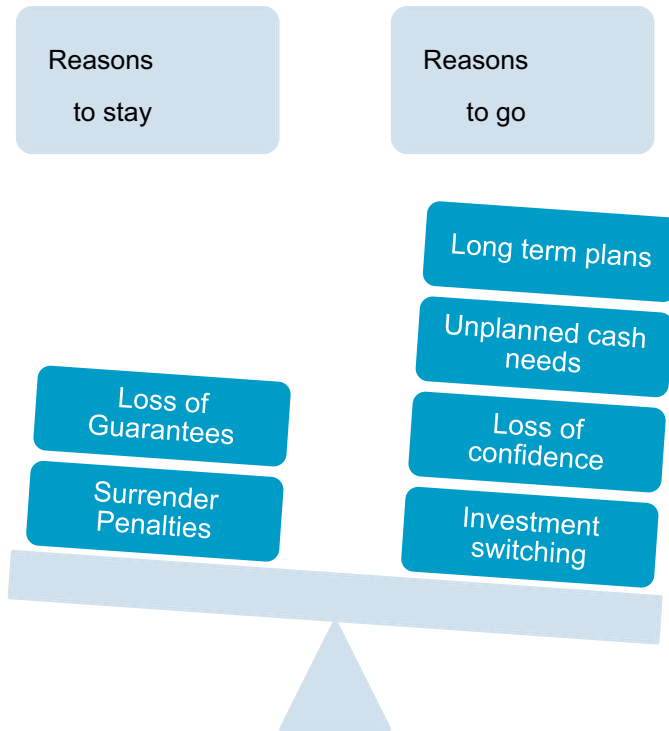
We are trying to get away from fitting probability distributions to data and think more about what is actually driving lapses. We can look at an example here with a savings policy as illustrated in Figure 30. The decision for every policyholder about whether to lapse their policy comes down to the weighing up of different options. There are reasons to go and there are reasons to stay. If the reasons to go start outweighing the reasons to stay, then policyholders are going to lapse. That is the way it works on an individual policy level. When you think about this for a fund full of policyholders, then, rather than it just being stay or go, it might translate to a lapse rate. If all the reasons to go start outweighing the reasons to stay, you might translate that to a particular lapse rate.

I will introduce the idea of a lapse factor function as shown in Figure 31. The lapse factor function is meant to represent the propensity to lapse. This is a function that turns the policyholder's desire to lapse into an actual lapse rate. We define it so that a lapse factor of 0 means no-one is lapsing. The higher the lapse factor, the more lapses there are going to be. Of course, the most there will ever be is 100%. You could think of this as the instantaneous lapse rate or the monthly lapse rate or something similar – something that you can use for a long timeframe. The idea is that, if we can model what affects this lapse factor, then we can calibrate the lapse factor function and therefore translate it into a lapse rate. We can now consider some of the different ideas that drive the lapse factor.

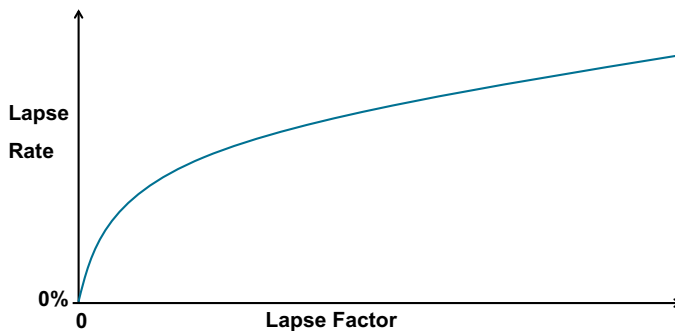
First of all, there might be long-term plans. This policyholder might have taken out his policy with the intention of lapsing it within a few years because he wants to get his money out to get on the housing ladder. You might not have very good information about the intentions of people. In that case, you should obtain better information. You could ask people, when they take out their policy, “do you have plans to lapse at a particular point?” Or ask them when they have lapsed. You need to try and get as much information as possible about the reasons for these lapses to be able to model them effectively.

The next lapse point is unplanned cash needs. For example, a policyholder's car has broken down and he needs to get a new one. That is an example of why he might need to get his savings out in a hurry. What kind of factors might drive that? Policy size might make a difference. For example, on really big policies, maybe unplanned cash needs would not usually drive lapses.

Lapse risk (savings policy example)



Lapse Factor function


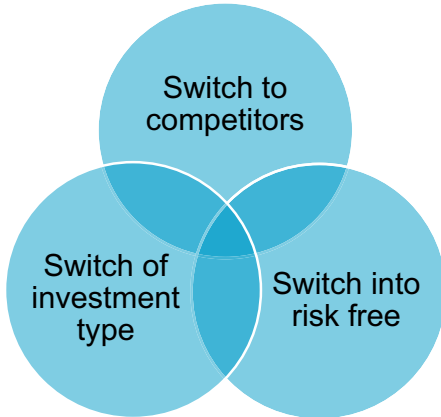


Economic conditions might be another factor. If there are really tough economic conditions, maybe people are losing their jobs more and need to access their savings more often.

The next point is all about reputational lapses. For example, if an operational risk event has just happened, that might increase the number of lapses. If your Solvency II surplus level is getting very low, maybe the reputation of the company is going downhill, and maybe people would start lapsing. Or maybe your Environmental, Social and Governance rating is getting poor. These are often the drivers of mass lapses.

Then there is the idea of investment switching. People might switch to a different type of investment or to a competitor. One possibility is that people switch to risk-free investments if market volatility is really high. In the example function shown in Figure 32, we are assuming low levels of switching at the beginning and end of the policy, but some more in the middle.

Investment switching



Next, we allow for things like surrender penalties and guarantee periods. These are certainly going to affect people’s desire to lapse. We need to model these and understand how significant an effect these things can have.

We take all of the things above and apply some established techniques like validation and back-testing to build a lapse model as shown in Figure 33. We then need to calibrate it. Excuse me for


Calibration

Calibration of lapse rate function based on:

- Causal lapse data
- Expert judgement

Supported by

- Backtesting
- Independent validation
- Documentation



slightly glossing over the details, but as I said, this is not meant to be a lapse presentation. It is meant to be about stochastic long-term modelling.

Many of you might be thinking, this sounds like it will be very difficult to do without a great deal more data, and it is going to be very reliant on judgment. The thing I would say to that is, great, let's get hold of all the data. This is the age of big data. This is the age where the likes of Amazon and Facebook have got reams and reams of data on us and are using that to great purpose, employing all sorts of data scientists. So let us get out there and get as much data on lapses as we possibly can and use that to model lapses as effectively as possible. We have had about ten years of very heavy work on proxy models, where many actuaries have had lots of fun using machine learning and other techniques for proxy model development. We should be using all these advanced data modelling techniques for lapse modelling to try and understand as well as possible what causes lapses.

It will certainly require some expert judgment to put these things together. But then, if you look at the existing framework we have, when we try and come up with a 1-in-200 lapse event, a huge amount of expert judgment goes into that as well. I do not think it is necessarily more challenging than the place where we currently are.

Figure 34 on lapse modelling is a comparison about what you can do with this kind of causal lapse modelling compared with a traditional approach. By a traditional approach, I mean fitting assumptions and distributions to data. That is what we would normally be doing. If we look at how we derive the Solvency II best estimate, the best estimate would be the probability-weighted best estimate of future cash flows. That is the Solvency II definition of what the best estimate is supposed to be. This has just been estimated straight away without any trouble, unlike under the traditional approach of fitting distributions to data. What firms seem to do is have a best estimate, and then make a possible allowance for non-linearity, which attempts to look at the different risks and therefore approximate the best estimate according to the Solvency II definition. Similarly, probability distributions can be obtained directly from the causal model, because we have stochastically modelled the risks. We are able to directly find the probability distribution of lapses. This could be over a one-year time horizon or over any other time horizon. Under our more traditional approach, what we tend to do is come up with the best estimate and a 1-in-200 stress, and then say it is normally distributed or distributed in a way similar to a normal distribution. That is not to say that lapses are all normally distributed or that the normal distribution is the best distribution to pick. Rather, the normal distribution becomes the default because that is the one people learned about in school.

Causal lapse modelling uses

	Traditional	Causal
Best Estimate	based on recent experience, (possible allowance for non-linearity)	Probability weighted best estimate of future cashflows
Risk Distribution	Distribution fitted to historical changes in lapses	Distribution fitted directly to lapse simulations
Relationship with other risks	Correlation assumptions used within Gaussian or T Copula	Causal relationships directly incorporated in the model
Flexibility	Applicable to the one year time frame within SII	Applicable to any time frame

The next point is relationship with other risks. The causal model is incorporating all the risks directly into the model. If there are causal relationships between lapse risk and equity risk, for

example, then you are directly allowing for these rather than just through statistical assumptions like correlations, and that can be a much stronger approach. The final point is flexibility. Causal models can be used for any timeframe you like. The distribution-fitting approaches tend to only work well for the one-year timeframe you have fitted them to. It is quite difficult to fit them to other timeframes.

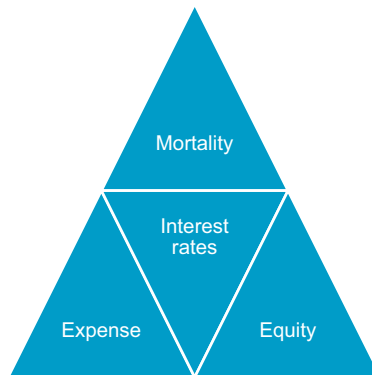
We will go on next to an example model an overview of which is shown in Figure 35. With some of the techniques I have talked about so far, it is really important to focus on leveraging existing processes, ESGs, stochastic mortality models, balance sheets and so on. Unfortunately, my example model is not sophisticated. This is just a simple model built with a spreadsheet and I was not able to use existing ESGs or stochastic mortality models or anything like that. It is, necessarily, a lot simpler than some of the things we have talked about. It is just to show what you can do on the strategy side with some of these models. This example considers some closed annuity business that is backed by some fixed-interest bonds, with the free assets in cash and equities. Therefore, we are exposed to mortality, interest rate, expense and equity risks.

Overview

Proprietary firm

- Administers closed annuity business
- Liabilities backed by gilts
- Free assets in cash and equities
- Uses SII SF
- No MA or VA

Risks



The objectives of the firm are as follows. We want to increase the value of the dividends. We want our dividends to be fairly stable and we want to avoid insolvency. Those are the three objectives of the company which are reasonable I am sure you will agree.

And we look at the strategy illustrated in Figure 36. On interest rates, the company is closely matching the liability with gilts, and rebalancing all the time to make sure that assets and liabilities are interest-rate-matched. That is, assuming they cover transaction costs. For equity risk, our equity backing ratio depends on the Capital Coverage Ratio (CCR). The idea is that we invest more in equities if we have got an additional amount of free assets. Similarly, with dividends, we have got a dividend strategy where we pay out dividends at a faster rate if we have a more healthy solvency position. This is the strategy. Later we are going to see how we can vary the strategy and test the aspects of it.

The risk model we are using is shown in Figure 37 and, as I said, it is quite simple. We are assuming that all the risks are independent rather than that there are any causal relationships between them. We use simple random-walk-type models. We have a single-factor Vasicek model

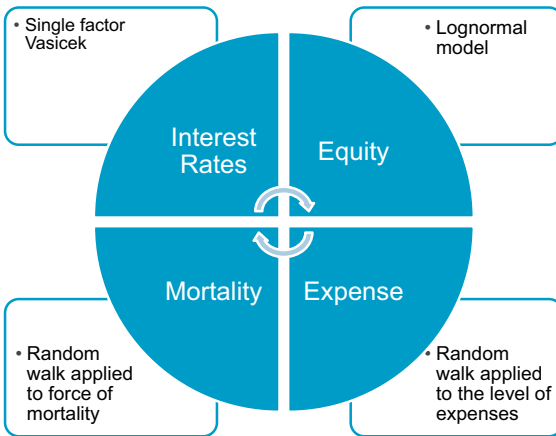
Strategy

- Interest rate hedging**
 - Liabilities backed with gilts
 - Regularly rebalanced to closely match
- Equity**
 - Equity backing ratio dependent on CCR
- Dividends**
 - Dividends linked to Own Funds, CCR

CCR%	Equity	Cash
>170%	50%	50%
<170%	20%	80%

CCR%	Dividend% of Own Funds
>200%	2%
160%-200%	1%
<160%	0%

Risk Model



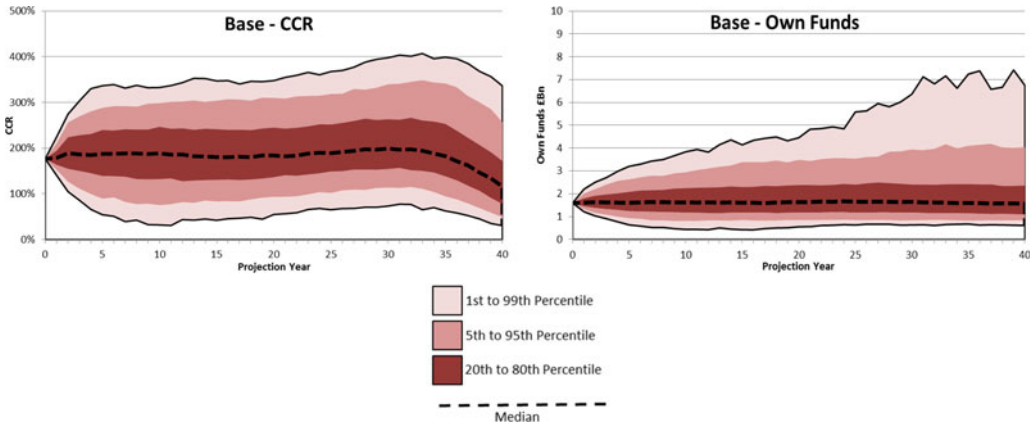
Risks assumed to be independent

for interest rate and a lognormal model for equity. It is necessarily a simple hypothetical example. In practice, you would use an ESG for modelling market risks and a stochastic model for mortality. We would use a causal expense risk model for modelling expenses.

The loss model for the annuity business is also quite simple. It is really just doing discounted cash flows in full. We are assuming the Solvency II Standard Formula applies here. The calculations just involve doing all of the runs in full. Best estimate liabilities are determined through a main run. The Solvency Capital Requirement (SCR) is calculated by applying standard formula stresses in full and recalibrating the stress value of the liabilities and assets. The risk margin is calculated using a risk drivers approach, probably familiar to those of you who use risk margins in practice.

So, that is the setup of the model. We will now look at some of the model results. First of all, we can look at the funnel graphs in Figure 38. We can use different shadings to represent different


Base Projections – Own Funds and CCR



percentiles and see how the business is going to look going forward. You can look at all sorts of metrics, but the ones I will focus on here are the CCR on the left and the own funds on the right. You can just see how these spread out over time. There are a few things to pick up on. First of all, on the left, the CCR tails off when you get past about 35 years. It is best not to worry too much about that because I have used a single model point here of 60-year-olds at the start. That means that they will be 95 or 100 years old towards the end of the graph. There is basically no one left by then, and you get some slightly silly results. With the own funds, there is a lot more upside potential than downside, and that is reflective of the risk mitigation actions we have. When the CCR reduces, we turn dividends off and reduce the equity backing ratio. That leads to a much bigger upside potential than downside potential.

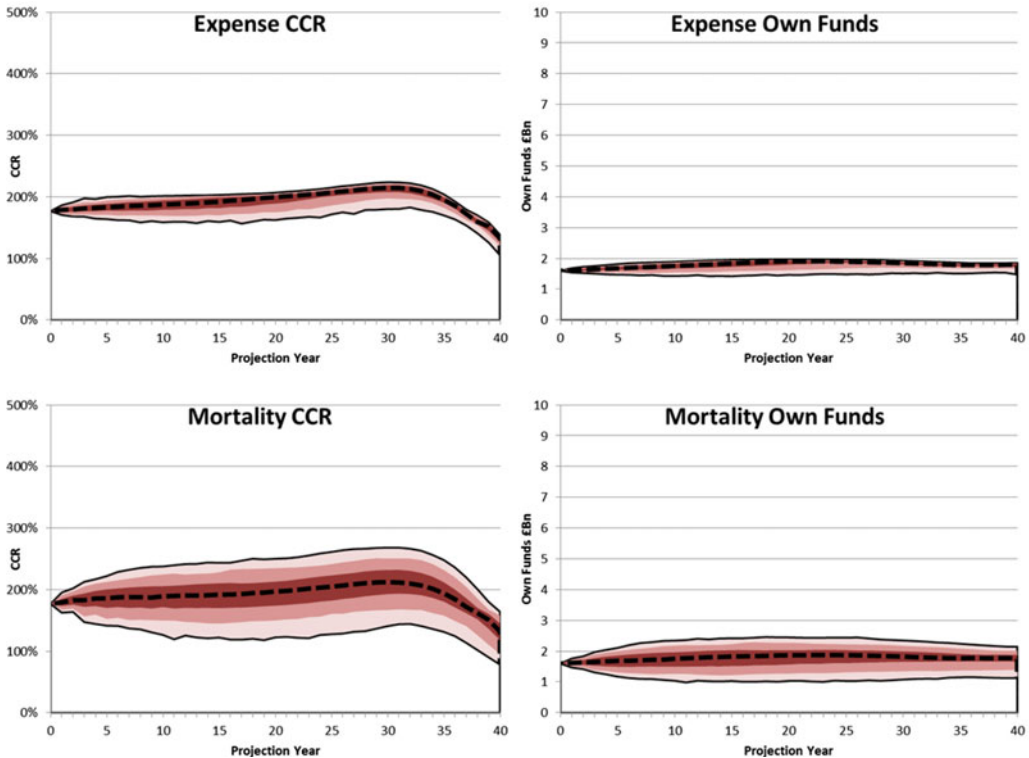
Next, we consider the surplus breaching probabilities, or the probability that the CCR falls below 100% as illustrated by the graphs in Figure 39. I think these graphs can be really helpful. They give the kind of insights you do not normally get. Many times I have thought about the question, as an actuary working in Risk Management, “What is the probability we can survive the next five or ten years?” It is generally very difficult to answer, but this class of long-term models can answer this question for you. You can look at the probability of breaches in any particular year, or the cumulative probability of breach. It is really only through stochastic modelling over a long term that you can get these insights. In this example, the probability of a surplus breach increases after about five years and then reduces after that. I think the reason for this is that we are more exposed to mortality risk earlier on.

Base Projections – Surplus Breach (< 100% CCR)



We can also look at individual risks as shown in Figure 40. Here, we are looking at expense risk and mortality risk. We can see that there is not much expense risk early on. There is much more mortality risk in this book of business than there is expense risk. The graphs are a little bit flat and boring, probably because I have used a single model point. In a real block of business, you would have policyholders of different ages. You might find that the mortality risk varies at different points over the projection. You would be able to see where your mortality risk and expense risk emerge over time, and what are the particular points to worry about.

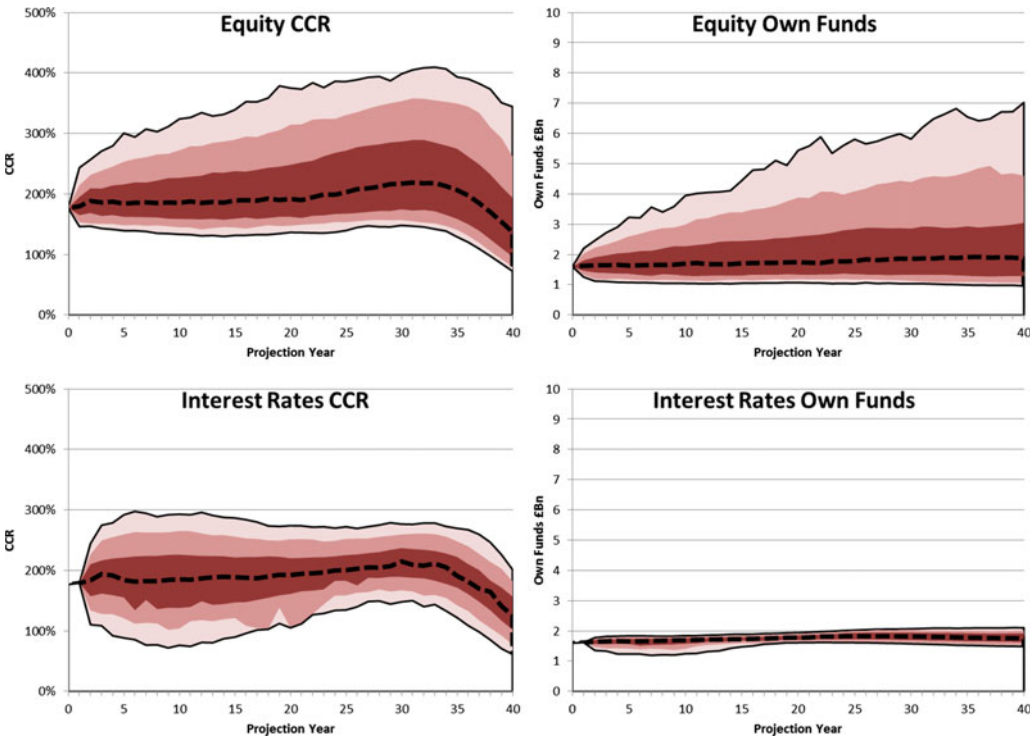
Individual risks 1



The same thing goes for equity risk and interest rate risk as illustrated in Figure 41. On equity risk, again, we have big upside potential, but not so much downside as a result of our risk mitigation strategies. Interest rates pose very little risk to the own funds, but do pose some risk for the CCR. That reflects the hedging strategy. We are hedging just the asset-liability position. The only variation in the own funds is where the risk margin varies. But, because we are not hedging the SCR, the CCR will fluctuate even if you are hedged on assets versus liabilities.

Now, we look at testing different strategies. We have three different strategies that the firm can consider here as shown in Figure 42. First of all, they might consider outsourcing the expense risk to another firm. As option two, they might consider just paying out the dividends a bit faster. Finally, they might consider a reinsurance strategy whereby they reinsure 90% of their longevity risk. We are going to look at these three different options and see how these fit in with our objectives of maximising dividends while keeping them quite stable and avoiding insolvency.

Individual risks 2



Strategy testing

Option 1

Outsourcing

- Admin company takes over expenses
- Outsource Premium reduces net assets
- Expense risk removed

Option 2

Increased Dividends

- 50% increase in assumed dividends

Option 3


Reinsurance

- Reinsurance premium reduces net assets
- 90% of longevity risk removed

CCR	Dividend % of Own Funds	
	Base	Option 2
>200%	2%	3%
160% to 200%	1%	1.5%
<160%	0%	0%

First of all, we can evaluate these strategies by looking at the average present value of the dividends as shown in Figure 43. As you might expect, option two, the extra dividend strategy, is pretty good for that. The outsource and reinsurance strategies are also boosting the dividends, but the extra dividends strategy is basically maximising this metric. But, as I said earlier, a lot

Strategy testing



of strategy testing is about balancing risk and reward, so you cannot just look at reward which, here, is the average present value of dividends. You also have to look at the amount of risk to decide on the appropriate strategy. To do this, we can look at the standard deviation of the present value of dividends. If we look at the graph on the right of the Figure, there we can see that although the extra dividends strategy is giving us a higher present value of dividends, it is also giving us substantial higher standard deviation as well, so maybe it is not quite so good as it might initially seem.

We can plot these two ideas, the average present value and the standard deviation, on the same set of axes in a graph like that shown in Figure 44. We really want to be moving towards the right and towards the bottom of the graph. That gives us higher dividends and lower volatility. We can see here that the outsourcing strategy does not have much effect, because we have not got all that much expense risk. The reinsurance strategy is increasing the value of dividends and reducing the standard deviation, so that looks like it could be a good idea. Then we have got the extra dividends strategy. That is increasing the value of dividends, but it is also increasing the standard deviation. So, that may or may not be a good strategy. We perhaps need to look at things in a slightly different way by considering some other measures.

Strategy testing - Dividends






Figure 45 shows our surplus breach probability. We can see how the different strategies alter this. First of all, the reinsurance one is represented by the red dotted line. If we go for this strategy, it almost eradicates the probability of a surplus breach until we have very little business left to project. That seems a good idea. The outsourcing strategy is reducing the surplus breach probability a little in the early years. It is only making a small difference, but it is quite helpful. Then we



have got the extra dividends strategy. It does not actually make that much difference. Even though we would be paying out all those extra dividends it is not actually increasing the probability of breaching surplus very much.


Next, we will look at probability distributions as shown in Figure 46. On the left, on the vertical axis, we have got the probability density. On the horizontal axis, we are looking at the value of the dividends. The thick, black line is our base assumption, and you can see what the different strategies do to the assumption. We really want to try and move this graph to the right as much as possible in order to get higher dividends. For stability, we would like a narrower, taller peak. We can immediately see that the reinsurance strategy does that very well. Also, if we look to the far left of the graph, it really cuts out the very worst outcomes. So, the reinsurance strategy looks good.




The outsource strategy does not make that much difference, but you can see that it is a bit lower than the black line on the far left. In other words, it reduces the risk of really negative outcomes. Lastly, the extra dividends strategy generally moves the line to the right to a fair extent, as you would expect because the average value of present value of dividends is higher. Looking on the far left of the graph, much like the outsource strategy, it has also reduced the downside risk quite a lot. It has actually reduced more of the downside risk than the outsource strategy. This shows that standard deviation is not the be-all-and-end-all when you look at the variation in dividends. Here, with the extra dividends strategy, we have increased the average value of the dividends

and increased the standard deviation, but standard deviation has only increased because of the higher upside potential. It has actually reduced the downside potential.

We can next talk about the utility functions as shown in the graph on the right-hand side of Figure 47. The idea of utility functions, for those who may not have looked at this for a while, is that they reflect investors' preferences. Investors are risk-averse, so you get concave-shaped utility functions that mean that investors really hate losses. Therefore, the graph is steeper on the far left and shallower on the right. We can apply utility functions to probability distributions to calculate the expected utility for each of different strategies and the results are shown in Figure 48. These show that the reinsurance strategy is looking like the best one and the extra dividends one the second-best strategy.



Expected Utility $E[U(x)] = \int_{-\infty}^{\infty} f(x)u(x)dx$
 $f(x)$ is the probability density function of the PV of dividends
 $U(x)$ is the utility function $U(x) = 1 - e^{-0.002x}$



That is the conclusion of those strategy exercises. We looked at different things and we have probably decided that reinsurance one is the one to go for. Another way of thinking about this strategy testing is that I have just tested three different strategies here, but in reality, you might think of these as continuous variables: the amount that you want to reinsure, the amount that you want to pay out in dividends, and the level that you want to set the equity backing ratio. With a

number of continuous variables like that, you can perform a strategy optimisation. You could derive the optimal values of these variables in order to produce the highest expected utility.

In summary, we have talked through some of the different generations of models and how it is often the rise in new technology that gives rise to a new generation of models. We have discussed some of the limitations in the existing one-year value-at-risk framework. We have further considered how we can therefore implement a stochastic long-term framework for modelling.

Some final thoughts. First of all, life insurance is really long-term business. So, is a one-year framework really suitable for it, or is something longer term better? Finally, we all have lots of fun doing modelling, but let us not get too carried away. We should always bear in mind that a model is just a model, and you need to have good insights into and understanding of the model. This includes things like communicating it. Understanding the assumptions and the limitations are very important aspects of the model, as is the judgement about how to use the model results. Do you have any questions?

The Moderator: That was a really interesting talk. I do remember the Wilkie model from years ago. One of my reflections on the SCR modelling we have been doing in more recent years is that there is a very sharp focus on the 1-in-200-year event and what is going to happen over the next year or two. When you think about things like reinsurance strategies and so on, you can often be quite short-term in your management actions. But I always had a fondness for the ultimate run out approach. Working out the losses in the long term, especially in the with-profits estate work, I think is really important. Some of the risks that are starting to emerge at the moment, things like climate risk or even the pandemic, just highlight the long-term effects of some of these issues.

When you consider this kind of risk, I often think of things like obesity. They just drip in over years and take a while to take effect. However, the assumptions you are modelling with today are not quite right for twenty years from now. You cannot really see that in the one-year model, but more so in that longer-term projection. Your interest rate model drew out the fact that interest rates were at 10%. You have to think about how much they could fall, given where they are today. When you look at how movements can happen, there are two types of movement. One is the slowly evolving movement, so maybe an obesity development, which might happen over 20 or 30 years. And then there can be other movements that can happen quite rapidly. For example, the pandemic over the last year or two has happened quite dramatically. So, some change in NHS policy, for example, where everyone is getting a new cure, could radically change people's life expectancy.

The following questions were raised by the audience.

Questioner: Is there a similarity between credit card default modelling and loss factor models?

Mr Curry: Credit card default modelling is not something I have ever been close to. I do not really have much detail about the way that it is done.

Questioner: How practical would it be for firms to move to a run-off capital approach, and should the Prudential Regulation Authority (PRA) be reviewing this as part of the Solvency II review?

Mr Curry: A lot of relevant things exist already. If firms can use the existing tools, such real-world ESGs and stochastic mortality models, I do not think it is a big step to build in additional items such as for things like natural risk and expense risk. That is something that is eminently doable for some of the larger firms. It is quite difficult, perhaps, for some of the smaller firms.

In terms of the regulatory capital environment and whether a run-off approach should be what you do going forward, I always bear in mind that there has been a huge amount of effort from the UK and elsewhere over many years working towards the value-at-risk framework. I do not think there is going to be a massive step change soon.

Having said that, I think there are a lot of advantages to run-off capital. As I said earlier, I think it is much more intuitive. It is probably much more aligned with what the policyholders want. When they give their money to a life insurance firm, they want security that the firm is going to be there for paying them back in 20 or 30 years. They do not really want a hypothetical transfer

value to another provider. I also think it is much less prone to pro-cyclicality. I always worry about this one-year evaluation framework. It is quite pro-cyclical. We work hard to put all these things in to stop it being pro-cyclical. But it is inherently quite a pro-cyclical process, so I do think there are quite a few advantages of looking at run-off capital. That may not necessarily mean a complete change in the regulatory regime, it could just be firms using this approach much more in their ORSA to understand their risks in a better way.

Questioner: This is a very interesting modelling approach. Do you believe it is possible to extend the model to fit in the current Solvency II framework to provide regulatory capital requirements?

Mr Curry: Yes, I do. This is about looking at projections of your company's surplus, and having a position that is stochastically projecting not just your assets and liabilities to get run-off capital, but the full balance sheet. In other words, projecting the SCR, risk margin and perhaps the transitional measures for technical provisions (TMTP). It can be quite challenging to model this. It might be necessary to make some approximations along the way in terms of, for example, using various run-off drivers to help in different areas. But I certainly believe that you can stochastically project the SCR and risk margin and therefore the full Solvency II balance sheet in this way. Some of the processes we have got set up to do these things are a lot better than they were five or ten years ago. I think there are certainly challenges, but yes, I very much believe we can model how the Solvency II balance sheet position moves going forward.

The Moderator: So, you project forward on the real-world basis, but at that future time point, you need to create your risk-neutral scenarios to do the valuation or some kind of proxy model to work out your guaranteed position.

Mr Curry: Yes, I think the difficulty is that we have the market-consistent valuation. It works well for time zero when we have a market for these things, but it is generally very difficult to project forward because what you project is the market value of your options five or ten years down the line. It is very difficult to understand these things. I think necessarily it has to be quite simplified, and based on the various run-off drivers, to try and get some insight into these things. We just have to understand that there will be approximations and limitations along the way.

Questioner: Should the capital modelling allow for climate risk and green finance?

Mr Curry: This is quite a new area of thinking in the actuarial profession. I am not sure we are quite developed enough to do it yet. I think that maybe in another two or three years, when we have done more climate risk modelling, we might understand this better. My initial thoughts are that we are probably best doing scenario testing at this stage. We could think of various different climate change scenarios that could happen, including how they might develop over time, and apply these as scenarios over your long-term projections. I think that is probably where we are at the moment. Maybe two or three years down the line we can think about stochastic modelling. Perhaps some firms are more advanced and they can think about doing it now. It is quite a challenge.

The Moderator: What would be really exciting is if you could map some of the projections about climate change that are out there in the public domain as some of your input drivers for your stochastic model. Then from there, you could model your asset prices, for example. They would be a model of what is happening in some of the underlying drivers that have been projected externally. You could have a really rich and interesting projection of how society could change over the next ten or twenty years. At the moment, if you look at the short-term horizon, I struggle to see exciting scenarios in terms of how things can develop. But if you could somehow connect the work that climatologists and geographers are doing on future risks to your model, that could be really fascinating in terms of giving some richer insights.

Mr Curry: Yes, I think so. I think it is important to build on things that exist elsewhere. You can do this by leveraging your ESG models and leveraging your stochastic longevity models, and that could certainly apply to climate risk models as well. There is a huge amount of expertise in these

specific areas and the idea is to bring these areas of expertise together in a single framework and be able to model things.

Questioner: Banks use logistic regression to model credit card payment default probability. So can you comment on whether you are using logistic regression in any of your loss factor modelling?

Mr Curry: No, I am not.

Questioner: Do you think that this type of approach would push actuaries to take stronger views on long-term economic drivers?

Mr Curry: I think so. I think we do not have enough focus on the long-term as a profession. I think we focus too much on short-term events. The capital framework we use is based on the concept that we are going happily along and then this big bang comes and hits us. I think it is much more important to think about the long-term drivers and look at much longer timeframes, because long-term risk is really important.

The Moderator: Risk projections are, of course, expected to be more challenging the longer the projection period. How far would you say we could reasonably project without the results becoming overly uncertain?

Mr Curry: I think it depends on the purpose. If you are just projecting the losses and therefore driving things like run-off capital, that is inherently a lot easier than projecting things like the full Solvency II balance sheet. If you are just projecting losses, I think you can reasonably do that to the end of each policy term and use that to project runoff capital and things like the runoff of with-profit estates. If, on the other hand, you are looking at full Solvency II balance sheets and understanding your surplus, it is probably very dependent on things like new business levels and the terms you might get for reinsurance, which are very difficult to model. Those things are so inherently uncertain. For that reason, I think you are probably best not going much beyond about five years.

The Moderator: What are the three most important things you would need to think about in the design of the long-term stochastic model?

Mr Curry: I think first of all, as I said previously, I am in favour of centralised design. You need to have a central point of contact that owns the design of everything. Things are then outsourced to the different areas and you leverage the different expertise, but you have it very much centralised in the overall design. Another important point is the governance of the models. We always talk about governance of the models under Solvency II. I do not find it an exciting subject, but it is so crucial. The reason it is key is not just that you have got your models under control so that you do not have changes happening at any stage. It is also about building understanding. It is really about having senior stakeholders understanding the way the model works and the limitations around it. That is almost as important a part of the governance as the actual control of the models themselves. I think governance is a really key issue. The last one is using what is available already. We should not try and reinvent the wheel but use things that are out there already. There are lots of brilliant ESG models and some brilliant stochastic longevity models. We should use them rather than build things from scratch.

The Moderator: We have a comment here from someone in the audience: These days, users are aware of the limitations of simple models. Making the model complex would better represent the risks, but there is also a risk that users are less careful and more reliant on the model. This is the black box dilemma.

Mr Curry: This is an interesting issue. These models can get quite complex, but it is similar to the situation with real-world ESGs. There is a huge amount of complexity in those. The number of actuaries who truly understand the way that these real-world ESGs work, to the extent they could build one from scratch, has got to be very small, but lots of people can still use them. Lots of senior stakeholders can look at results from them and understand them. It is a case of understanding what the model is trying to achieve, understanding the approximations within it, the assumptions within it and the limitations in there. If you can do that, you do not need to understand every last

bit of the detail. That is the case for long-term modelling in general. It is a case of not trying to understand every last technical detail. It is just about understanding the overall principles of the way these things fit together, and being very clear about the assumptions and limitations.

The Moderator: The Solvency II framework puts some focus on model validation for internal models. My experience is that model validation teams are generally opposed to expert judgment. This approach is deemed to increase the dependency on expert judgment. How would you tackle the challenge for model validation teams? Do you need more data or perhaps more academic research to support this approach?

Mr Curry: Well, I certainly think more data is great. I really think that, as I discussed with respect to lapse risk earlier, this is a huge growth area. We should be obtaining every last bit of data we possibly can to feed these models because data gives understanding. We need to really get as much lapse data as we possibly can and really understand these risks. There certainly will need to be some expert judgement, but there already is extensive use of expert judgement. In Solvency II, we talk about coming up with a 1-in-200 lapse event when we have, perhaps, ten years of data if we are lucky. There is a huge amount of expert judgment in there already. I would not necessarily say there is going to be more reliance on expert judgment.

The Moderator: I actually think there is a lot of data out there. I think you made a really good point that there is so much data that is untapped. The other way I look at it is that you need to consider what you need to believe. You build the model or the scenario and, in that world, you need to accept that there are certain constructs within your model. The problem with the over-simplified approach is you are essentially ignoring some solid data out there that might give you some richer insights. I think if you can connect your model to data, that is really powerful.

Questioner: Is there a risk of building models that get so complex that they may be poorly understood and regarded simply as a black box? This is similar to the question we had earlier, but with a focus on complexity.

Mr Curry: I think that risk is always present. Initially, I think the way to get around it is to make sure we have good governance so that people are fully aware of what is going into these models and what we are building. We cannot just have a process where actuaries just go away and start building things without any particular oversight, so no one therefore knows what is going on in the model. I think governance is key to stopping it being a black box. Also, setting out the assumptions and limitations will prevent it becoming a black box to some extent. I would also say that a lot of this approach is more intuitive than some of the existing models. Run-off capital is quite an intuitive concept: just having enough money to pay your liabilities with a particular probability. Hopefully, some of this takes you away from a black box and into some more intuitive areas. Perhaps the same applies with lapse modelling. You look at the causal drives of lapses in the real world. That is the kind of thing that people can understand. It is perhaps easier to get into some of that than it is to understand fitting fancy probability distributions to data.

The Moderator: On the causal side if you can connect with what your sales managers and others are looking at, and reflect that as factors in your model, it helps build buy-in for your model. It gives insights back that are consistent with the behaviours that they are anticipating. Hopefully, it gives them some new dynamics and features so that they can model and run the business a bit better.

Questioner: If we look to build causal relationships between all the risks in the model, does this then turn into an enormous exercise we cannot ever hope to complete?

Mr Curry: Probably the answer would be yes if you did try and build causal relationships for every single pair of risks in your model. It is a bit like correlations. If you have got a correlation matrix, you might have, say, 500 correlations to populate inside it. That also sounds like an enormous task. You would probably have to have an epic expert judgment session to get through them. But a lot of risks do not have much of a relationship. The thing to do is, rather than go through every pair of risks and work out what to do with them, focus on the real key areas where we expect there to be a relationship and concentrate on getting those right. These could be things like, say,

equity markets or operational risk events driving lapses, or falling interest rates driving GAO take-up.

The Moderator: Thanks so much, Bill (Curry), for a really interesting talk. Also thanks to everyone for joining today. I hope everyone really enjoyed the session.