

---

# Pension Plan Solvency and Extreme Market Movements: A Regime Switching Approach

---

## Abstract of the Leeds discussion

[Institute and Faculty of Actuaries, 11 June 2012]

### Contact

Mark Freeman, E-mail: M.C.Freeman@lboro.ac.uk

This abstract relates to the following funding report: Abourashchi, N. Clacher, I., Hillier, D., Freeman, M., Kemp, M. & Zhang, Q. Pension Plan Solvency and Extreme Market Movements: A Regime Switching Approach – Funding Report for the Actuarial Profession. *British Actuarial Journal*, doi: 10.1017/S1357321713000287

**The Chair (Mr M. Slee, FIA):** Thank you for coming to tonight's discussion on pension plan solvency and extreme market movements: a regime switching approach. I am secretary of the Yorkshire Actuarial Society.

I will now hand over to my co-chair, Gareth Connolly, of the Pensions Practice Executive Committee, to chair the discussion part of this meeting.

**Mr G. T. Connolly, FIA (co-chair):** Thank you for inviting me and thanks to the Yorkshire Actuarial Society for hosting this session. I am on the Pensions Practice Executive Committee. I will ask Dr Clacher and Professor Freeman to present their paper.

**Prof M. C. Freeman:** I should like to thank you very much for inviting us. I am from Loughborough University; Dr Clacher is from Leeds and we are two of the authors of this paper.


We have received financial assistance for this research from the Rotman International Centre for Pension Management, which is based at the University of Toronto, and from the Actuarial Profession in the United Kingdom. We would like to thank both for their support.

What we are going to look at tonight is the future funding position of defined benefit (DB) pension plans. The funding position of the plan, under a mark-to-market convention, is the difference between the current market asset value and the present value of the liabilities as a fraction of the present value of the liabilities.

You may be familiar with this chart in Figure 1. It is from the Pension Protection Fund and shows the aggregated funding position of the UK DB pension plans that it covers. Below the horizontal axis is underfunded; above is overfunded. You can see that this line is highly volatile, going from 20% in surplus in mid-2007 down to 20% in deficit two years later. Since then, it returned to being slightly in surplus before dipping back down into a deficit again of around 17% or 18% at present.

## Aggregate funding position of UK DB pension plans

The Pension Protection Fund 7800 Index. Source: <http://www.pensionprotectionfund.org.uk/Pages/PPF7800Index.aspx>



Mark Freeman (Loughborough University)

Pension Fund Solvency

IoA, June 2012

4 / 35

**Figure 1.** Aggregate funding position of UK DB pension plans

The question that our team set ourselves was: “Can we project this line forward, either at the aggregated level to help the Pension Protection Fund and other macro-prudential regulators, or at the individual fund level to help pension funds, to estimate the future risk of underfunding?”

What influences this chart? Some of the issues are fairly obvious: changes in life expectancy and asset performance, for example. The issue that a lot of our discussion today will focus on, by contrast, is the choice of discount rate.

What you are looking at is not the liabilities themselves but the present value of the liabilities, which will depend on the discount rate used in order to calculate the present value.

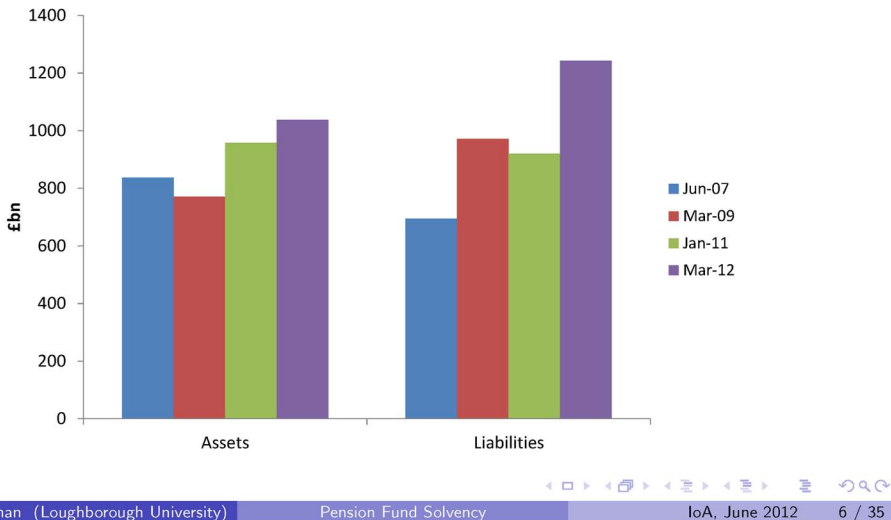
So let us go back to the above chart and examine what caused all the movement. On the left of this next graph in figure 2 is the asset value, while on the right is the present value of liabilities.

What you can see is that the big drop in solvency between 2007 and 2009 was not predominantly caused by asset values dropping. While the equity market fell, the bond market rose, and these two effects partly offset each other. So, in aggregate, there was not a huge movement on the asset side of the balance sheet.

By contrast, there was a significant movement on the liabilities side of the balance sheet, not so much because the liabilities themselves were changing but because bond yields were decreasing, resulting in a substantial rise in the present value of liabilities.

## Assets & liabilities of UK defined-benefit pension plans

The Pension Protection Fund 7800 Index. Source: <http://www.pensionprotectionfund.org.uk/Pages/PPF7800Index.aspx>



**Figure 2.** Assets & liabilities of UK defined-benefit pension plans

If you recall, between 2011 and 2012 pension funds in aggregate went from marginal surplus to large deficit. This happened at a time when asset values increased. But the present value of liabilities rose more steeply.

Therefore our first contention is: if you are going undertake time series modelling in this field, you need to put more attention on the liability side of the balance sheet than you do on the asset side.

Who should be interested in our findings? At the individual fund levels, managers, actuaries and trustees should all find this work relevant. But we are also interested at the aggregated level, and particularly how the Pension Protection Fund might use this type of analysis to work out the insurance premia that they charge to pension funds.

In addition, our work will have an impact on industrial relations, particularly in respect of the amount of contribution that people are going to have to make to DB funds and the types of benefits that they can expect to receive. All the current arguments concerning contribution rates and final salary terms will also be tied into our assessment of future pension fund solvency risk.

Our methodology is based upon a frequentist approach that relies heavily upon historical data. Half of my academic career has been spent talking about how the past is not a very good indicator of the future. But, for the purposes of tonight, given that we have a limited research question, we are going to abstract away from the limitations of using past data to project future performance.

Given this frequentist approach, our specific research question is: “What is the probability that a given DB scheme, or DB schemes in aggregate, will be less than 100% funded at some future point in time?”

To address this question, our econometric techniques will capture two important characteristics of the data. The first is its ability to model fat tails in asset returns, which is consistent with observed non-normality. The second is that we are going to jointly model the asset and liabilities sides of the balance sheet.

By contrast, what we are not going to focus upon is the specific liability cash flows. Instead we will model a stream of liabilities. Our model is not restricted to such a scenario, but it is appropriate for any liabilities structure that you may wish to consider. However, for presentational purposes we do not want to over-complicate matters. In addition, we are also not going to look in detail at the inflation process which drives the nominal liability process. Again, our model could easily be adapted to incorporate such effects.

We have six models to present to you tonight. The first is the simplest model, where the discount rate is fixed and asset returns are jointly lognormal.

In the second model, the risk-free rate process follows an Ornstein-Uhlenbeck process, and asset returns remain lognormally distributed and independent from changes in the risk-free rate.

The third model is similar to the second model but now we allow for the observed correlation between the discount rate and asset returns processes. This is crucial. As yields go down and bond prices go up, you would expect there to be a negative correlation between the asset and discount rate processes. This is captured in the third model, but not the second.

Models 4 to 6 mimic models 1 to 3, but we now capture a Markov-switching environment, which introduces leptokurtosis into our simulations.

Our central conclusions are as follows.

First, we find that estimates of future pension fund solvency risk are highly model-dependent.

Second, it is essential to account for stochasticity in the discount rate process because much of the volatility in DB funding positions comes from the present value of liabilities.

Third, allowing for fat tails in the asset return process significantly increases pension fund solvency risk estimates.

Fourth, ignoring the correlation between asset return and discount rate processes leads to a significant overestimate of funding risks.

We present the results for just one illustrative portfolio which looks similar to a standard UK pension fund: 25% in UK equities, 20% US equities, 40% in gilts, 15% elsewhere internationally. All estimates are calculated on a UK sterling basis. Discounting is undertaken using a Treasury bond yield, rather than an AA corporate bond rate, but that is not going to make a substantive difference to the results.

Our illustration assumes that the DB plan is currently 15% overfunded. Again, it is easy to change this parameter, but our baseline scenario reflects the position of funds in the middle of 2007, before the current funding crisis occurred.

The liabilities are highly stylised. Effectively, they are nothing more than a 30-year growth annuity with no stochasticity in the liabilities themselves.

It is not our contention that this represents a realistic liability structure, but our focus is not on modelling the liability cash flows. Our general framework could easily be applied to a more sophisticated cash flow process.

We choose the initial cash flow so that the fund is 15% solvent in Year 0. We fix the inflation rate at which the liabilities grow at 4% on an annualised basis.

In this case, because it is a standard growth annuity, there is a simple formula for the solvency of the pension fund, which depends partly on the market value of the asset and partly on the growth annuity value, which in turn depends on the discount rate.

The movement in the present value of the liabilities is not coming from the movement in the liability cash flows themselves, which are effectively fixed; instead, it is coming from the movement in the discount factor.

Model 1 is the simplest one with a fixed discount rate equal to 3.6% which was the market value in December 2010 when we finished our calibrations.

For the asset side of the balance sheet, we assume that logarithmic returns are normally identically and independently distributed (n.i.i.d). This results in expected returns and the variance/covariance matrix between different asset classes being constant over time. We calibrate our model using monthly data from February 1970 to December 2010.

What we find in terms of projecting forward solvency risk – remember that the fund started 15% overfunded in Year 0 – is that in approximately 24 months, there is a 4% probability of your DB scheme becoming underfunded. The risk of underfunding then drops off dramatically. The reason for that is that the inflation rate of 4% is a long way below the expected return on the portfolio. In the long run, the portfolio returns dominate the inflation effect on the liabilities.

The second model incorporates a standard Ornstein-Uhlenbeck, or AR(1), process for the risk-free rate. From this we transform the risk-free rate into the variable  $x_t$ , which is also n.i.i.d. Transforming the variable of interest from an AR(1) process to an AR(0) process makes the analysis easier.

So, for Model 2 we still have the standard one state process for asset returns but we have introduced stochastic discount rates.

Again, you can work out fairly simply, using numerical integration, what the probability is of being underfunded at any given point in the future. In contrast to Model 1, we now find that the maximum probability of becoming underfunded is 17%.

Remember most of the volatility in funding positions comes from the liabilities side of the balance sheet. It makes a huge difference whether you treat the risk-free rate as being a fixed variable or whether you treat it as being stochastic. That is our first main result.

Model 3 is effectively the same as Model 2, but this time we capture the fact that the discount rate tends to go down when bond and equity prices are going up. This leads to a negative correlation between the discount rate and the portfolio value.

If you account for this co-movement, you take out some of the stochasticity in the system. Once we account for the fact that these tend to move together, this brings down the probability of default. When the asset values are going down, the present value of liabilities generally are going down as well.

You can see the maximum probability of becoming underfunded is only about 13%, compared to 17% in Model 2, but still considerably more than it was under the first model.

Now this finding is, in the context of current market conditions, somewhat counter-intuitive. What we have seen at times in the current credit crisis is interest rates and asset prices going down simultaneously. This, though, is unusual. Over the data that we have analysed going back to 1970, these variables generally move in opposite directions.

To summarise, under the basic model, solvency risk does not look very severe. But once you start to allow for the correlation structure between the interest rate and the asset returns process, this suddenly starts to become much more serious. Under what I would call our best model so far, there is a 13% risk of moving from a position of 15% overfunded to becoming underfunded within about a two-year period.

The problem with modelling single state normally distributed returns is that you do not get the fat tails that we all know are present in the data. One way of capturing this leptokurtosis – it is not the only way – is to use what are called Markov-switching models.

In a Markov-switching model, there are different states of the world and, within each state, returns are normally distributed. But the distribution of returns is different in each of the different states. By moving from one set of normally distributed returns to another, you start to generate leptokurtosis in the unconditional distribution.

In our model we have four states of the world and in each of those states the expected return vector and the variance/covariance matrix are different. They can be pictured as a bull state, an average state, a bear state and a crash state. There is also a fixed transition probability matrix that gives the probability of moving from one state to another.

I will not go through the literature in detail, but this type of model has been applied quite extensively to related problems, particularly value-at-risk.

This type of modelling brings three advantages: it incorporates fat tails, it allows for time-varying variance/covariance matrices, and it incorporates crash market states, which is one of the things that we were initially interested in capturing.

There are, as always, disadvantages. First of all, it is not easy to estimate and it is necessary to use some reasonably sophisticated software. There is also a large number of parameters. In our example, there are five asset classes and four states, resulting in 92 parameter values. Because of this, we often find, as is common with numerical estimation, that the software does not always converge without some transformation of the data.

I will go through the four states. The first one is a low returns state, with expected returns of about 0.7% a month. The standard deviation is fairly normal and the ergodic probability, which is the proportion of the time that the model spends in this state, is about 41%.

The second state is the bull state with almost double the expected returns of the previous state. Again, this occurs about four months out of 10. This state has below-average volatility.

Then we have the two slightly rarer states. State 3 occurs 15% of the time and is characterised by its low volatility.

The final one, which occurs 7% of the time, is the crash state with annualised returns of about -15% or -16%. It also has a standard deviation two and a half times higher than the normal.

We first simulate portfolio returns and estimate the first four moments and compare these against the data. The key thing here is that we are capturing the kurtosis in the distribution.

Model 4, like Model 1, has a fixed discount rate but now uses the Markov-switching process for asset returns. You can see, by comparing the results of Model 1 and Model 4, that incorporating fat tails has more or less doubled the perceived risk of becoming insolvent. This has risen from a maximum value of about 4% to about 7%.

Model 5, like Model 2, captures a stochastic discount rate that is independent of the asset returns process. This gives us the highest risk of all at approximately 20%. What is also apparent here is that the risk dies off more slowly than in Model 2. So funds face a longer-lasting risk because the fat tail properties stay in the data for longer.

Model 6 incorporates the discount rate process into the Markov-switching model to capture the correlation between interest rates and asset returns in the different states of the world. To do this, we include  $x_t$ , which is also an AR(0) process, as an extra variable to be estimated within the Markov-switching model.

Because we are now adding a sixth variable, we need to re-estimate the Markov chain. Now the most interesting state, with ergodic probability of 6%, is again a crash state with very low expected returns and very high standard deviation.

Within these crash states, the correlation between the discount rate and the asset returns process becomes less negative but still remains negative. Remember, in the recent crash it seemed as if asset returns and the discount rate were moving in the same direction. In general, though, in crash states they do not. In all four states, the asset return process and the discount rate process remain negatively correlated.

If we look at the unconditional distribution, we find that we do slightly less well in capturing the excess kurtosis. Once we start to include the interest rate process in the estimation, we do not capture the fat tails quite as well, although this remains substantially better than in Models 1 to 3.

And so we come to our final model, which we think is the best. This captures fat tails in the asset returns process and a correlated Markov-switching Ornstein-Uhlenbeck discount rate process.

Here the maximum probability of default is about 15% with the solvency risk again declining very slowly.

So what are our conclusions? If you are trying to model the risk of DB scheme underfunding, the result you are going to derive will be highly model-dependent. In our case, the spread is from 4% maximum risk for the simplest model to 20% maximum risk for Model 5.

To do this well, there are three key features that you must incorporate in your model: allowing for stochasticity in the discount rate, allowing for tail risk in the asset returns process, and co-estimating those two features of the data.

**Mr Connolly:** The modelling itself is, by necessity, quite simplistic, as it assumes the schemes start with a surplus, and it looks at the pension fund in isolation so does not allow for new money from anywhere else.

Is it easy enough to extend this model to reflect reality, as most schemes are in deficit?

**Prof Freeman:** In terms of making the cash flow liabilities more realistic, we were struggling with an issue.

If you start to make the liabilities very complicated, then you add more degrees of freedom in your model and you do not quite know what is causing the effect.

Therefore, for our academic purposes, having simplistic cash flows makes it clear what is driving the final results.

However, if you want to make the liabilities more complex, it is straightforward to do so within this framework.

**Dr Clacher:** One of the reasons we kept the liability model simple is because we wanted debates about models of future solvency to be the focus. If we start to have overly complex models of liabilities, we end up debating about liability modelling rather than the focus of the paper.

**Mr R. A. Leake, FIA:** You mentioned a number of assumptions you made about variables being independently and normally distributed, and so on, fixing one or two of them because it was too difficult not to.

To what extent is the model robust against the moves that you have had to make in order to make progress?

**Prof Freeman:** Dr Clacher and I were talking about this on our way here. It is to do with parsimony. Whenever you model anything, you want it to be complicated enough to capture sufficient complexity to make it interesting in a real-world context. At the same time you want to make it soluble.

I think the interesting issue for us, having seen these conclusions, is that you need to go at least as far as Model 6 because the results are so model-dependent that even at Model 6 you are probably not at the end of where you want to go. But nobody to our knowledge has even reached Model 6.



So does the model need to go on and become more complex? Yes, I think it probably does. I think that we have demonstrated that you need it to be at least this complicated, rather than survive with something a bit simpler, like a lognormal model for assets.

**Dr Clacher:** It is also achievable if it is what you are using it for. If you want to get some feel for extreme movements and how they will affect pension funds in aggregate you need a sophisticated model.

However, the more complicated things become, the more the modelling itself becomes the end game. You lose sight of the problem again. While the basic model clearly has some merit, as we have different, more sophisticated models, we start to capture more things.

The data would help the decision-making around model choice and, at a certain point, would then go into diminishing marginal returns. We would like to say that the level of complexity in our modelling, in our opinion, is around the optimal point: eight states and 22 different asset classes. It starts to become computationally intractable. It is a real trade-off between a better model and better decision-making.

**Mr Leake:** It is a little bit like asset liability modelling. It does not tell clients the answer but it hopefully gets them to understand a little bit more about the level of variability there is and how things work.

As an educational tool it is useful, so long as you do not think this gives them definitive answers.

**Dr Clacher:** Any output from a model should be up for debate. The way we do it means you can present the output in different ways. You could have further states with rates of inflation for the liabilities. So the liabilities increase at a much higher rate than we have in the model and the results could be presented as fan charts. You can therefore show trustees or sponsors a range of outcomes. So it is a useful decision-making tool.

**Prof Freeman:** One of the concerns that we had when contemplating pension fund holidays back in the mid-2000s was that people had not realised how quickly it was possible to go from being overfunded to being underfunded. Therefore, we suspect historically they were making bad decisions about pension fund holidays on the basis of using the wrong model. It does not mean we have got the right one, but I suspect that we have a better model than they had.

**Mr Leake:** A lot of trustees and sponsors would no doubt prefer you to demonstrate a potential upside in your model of 30% and give them some hope that one day they might emerge from their current funding difficulties.

**Dr Clacher:** That is easily done. You can use it either way. You could decide you want to impose Japanese style stagflation. The model is flexible enough as a technique that you can bring in different views of the world and start to do quite complicated scenario analysis, such as assuming that asset returns are going to be terrible for the next 10 years. You could impose that on your model if you wished and you would then have a very different outcome. So it is quite useful at the policy level as well. It will not necessarily give the right answer, but you can start to look at Armageddon scenarios, for example.

**Mr J. V. Betts, FIA:** This detail and modelling are new concepts to me and I am intrigued by the idea of flipping between states.

I recall listening to P. D. Jones (who also published a relevant article<sup>1</sup> on equity and gilt returns) who considered that the gilt market over the last hundred years appeared to flip between two states, one of which was a low yield state and the other a high yield state.

I would suspect this could be correlated with the inflation issue that you have not yet taken into account.

**Prof Freeman:** One of my main worries about this is that we have stochastic discount rates and non-stochastic inflation rates. The option is to model the inflation rate too. But the inflation rate is generally fairly short and therefore the discount rate you use is the discount rate today; but if you are going to use an inflation rate over the next 30 years, you have to project it over the next 30 years. I am not quite sure how best to do that. There is clearly an issue with capturing the co-movement of inflation and discount rates, but the way in which you would do this in a model is not simple.

**Mr Betts:** How dependent on your approach or assumptions was the actual split into the four states that you have derived and then applied to the data over a period? It occurred to me that something quite useful for people who are trying to manage closed pension funds – which are typically in a negative position – is what we would do if we had more asset growth to reduce risk. When looking at the four states and one of them is the base state, which applies 40% of the time, could it be that part of that state is in fact a positive environment. In other words, that there is a fat tail on the upside as well as on the downside, albeit with a relatively small probability?

Has the analysis you have done led to you understanding whether that was the case or not?

**Prof Freeman:** The reason we have used four states is that there is a seminal paper by Guidolin and Timmermann, who looked at co-estimating stock and bond returns. They conclude that you need four states in order to do this. So, you certainly cannot get away with fewer than four states.

Could you add more than four states? You could, but then that increases the estimation problem. You start to get more uncertainty in the variables you estimate because you have to estimate more of them. So again, back to Dr Clacher's point, it is a trade-off.

**Mr Betts:** It does demonstrate the concept and, hopefully, although I am not a person to judge, will give some useful outcomes.

**Mr Slee:** Can I just follow on from those comments, bearing in mind I have a life background rather than a pensions background? One of the areas I am currently involved in is Solvency II. Part of Solvency II is demonstrating that senior managers understand their model; that they know how it works. It is integral to running the business and, therefore, they make business decisions based on the results.

Clearly, you are going to have exactly the same issue here. If you are saying you ought to be using these models, how do you get trustees and the companies to understand them? Is that going to be achievable for models like this?

**Dr Clacher:** There has to be effective communication of the output of the model. And there has to be an element of trust in what actuaries do, whether they are pensions actuaries or

<sup>1</sup> JIA 120 II 253-309.

insurance actuaries. You are not going to be able to get complex models fully understood in most circumstances.

Producing fan charts from the models will help the decision-making. But there needs to be trust and understanding between decision makers and advisors in this area.

**Mr Slee:** I accept what you are saying. Coming back to Solvency II, it may be unrealistic to expect senior managers to understand all the in-depth complexities of the model, but they should, as a minimum, be able to understand the concept and intentions of it.

**Dr Clacher:** This approach is a variation of value-at-risk, which is generally understood.

**Prof Freeman:** It is. I have had that comment in the past. I think the view that was expressed was that Markov-switching was an easier way than copulas. People have some concept of what the states are, so you show the four states and say there is a crash phase and there is a bull phase. You don't have to explain the technical details. Then there is a reasonably conceptual way of presenting this. My answer would be that it depends on the level of detail.

**Mr A. D. Buxton, FIA:** This does not appear more complicated than the work that is being done on the insurance side. I think if it is possible to explain things in a satisfactory way on the insurance side, then it should be possible to explain this.

Do you have a version of the graph illustrating the results for a portfolio that is more heavily invested in gilts?

**Prof Freeman:** No, we have only done one, but it would be easy to do. If you would like us to do it, send us the portfolio. We can do it in half a day.

**Dr Clacher:** In our model we have taken historical returns on gilts. It assumes the same historical gilts returns going forward. It is not assuming anything about gilts as an investment class and their suitability for pension fund investment. In terms of being underfunded, the outcome would be much higher than for a gilts portfolio because the returns for the gilts portfolio are going to be much lower over the future projection. This is not necessarily bad in terms of the modelling because many people are going that way.

**Mr B. S. Nelson, FIA:** On pages 10 and 11 of the paper, you set out the different states. I am intrigued that you have got the bear state as the most common state. The second most common is the bull state. That is slightly worrying. Most of the time, equities will underperform expectations.

**Prof Freeman:** I presented the paper at Sheffield a couple of weeks ago and they took me to task for using the term "bear state" as it still has positive expected returns. Therefore, in this presentation I changed it to "low returns state".

**Mr Nelson:** That is more common than high returns. If you looked at the data at the end of 2000, you would have different conclusions.

**Prof Freeman:** That is always the problem with the frequentist approach. Again, when I presented this academically, many people picked the model and said "Why have you chosen Markov-switching between four states?"

To me the key assumption is the way we calibrated it. We chose data from 1970 to 2010. Had we instead chosen 1930 to 2010, or more recently, then we would probably have got very different results. That is always the problem.

**Mr Connolly:** Can I make an observation and ask a question? The technical actuarial standard for modelling, or TAS M as we know it, basically says that you need to confirm to users of actuarial advice that the model is fit for purpose and explain any shortcomings, simplifications etc.

The observation is that, while models must be more precise and capture real-world events, trying to explain them to users of actuarial advice in a way that is consistent with the actuarial standards might be a battle in itself. Has anyone any thoughts on that?

**Mr K. L. Ternent, FIA:** You talked, Mr Connolly, about fitness for purpose. I have been looking at this from the other direction. The first thing the analysis told me was that it could be interesting to the Pension Protection Fund, and the service it performs, particularly as extreme events seem to be happening over the very short term, and that is what the PPF concentrates on.

So if we discount the PPF as a logical choice from a client employer's point of view, because it does not want to go bust, we have two other choices. We can either play a long game and aim to cover the buy-out liabilities, in which case we are into the insurance company actuary's nightmare of Solvency II, and few people have the prospect of a affording buy-out any time soon these days; or we have the middle prospect, which is "I cannot afford buy-out; I do not intend to go bust. What am I going to do that lies in the middle between these extremes?"

I recall a paper written by Daykin and others on discount rates, in which one of the early suggestions was that a discount rate is not useful if it produces volatile results. Any discount rate related to AA corporate bonds, gilts or any other rigid formula, over the past few years will appear to produce volatile results.

Given that, over the long term, provided we have a functioning economy, equity investment may be expected to produce higher returns than bonds, then, maybe we should just stick with equities and sit back for the next 60 years and watch those better returns happen.

**Prof Freeman:** A 100% equity portfolio, in my opinion, gives much lower returns in the future than it has in the past.

**Dr Clacher:** I think that is a great idea but I received criticism in the Financial Times for suggesting that. I have given up on that as a public standpoint for now as I keep being berated for it.

**Dr Clacher:** I think it is a question of balance with investment decision making. I have done some work with Marcus Hurd and John Hatchett on flight paths and setting targets for achieving pension outcomes that both trustees and managers agree on. Part of this is looking at the current investment strategy, agreeing where the scheme wants to be at a certain point in the future, and then agreeing a strategy for getting there.

**Dr Clacher:** What we actually said is that there are only a few decisions that matter, and the choice of discount rate is not one of them. It is a means to an end. There is far too much emphasis placed on this. Personally, I do not agree with risk-free discounting although we use it in

the current paper. All it does for us here, and we stated this explicitly in the paper, is allow us to have sufficient data for an interest rate series for discounting, because AA bond data for the UK is not sufficient.

I am an advocate of using the long-term expected return on the investment portfolio as the discount rate – back to a SSAP 24 type long-term actuarial view. But, at the same time, that in itself had problems because it under-estimated what was happening in the short term as it was always looking at the long term. Again, it is about balance.

**Mr Ternent:** It is a flexible approach to setting discount rates that I am suggesting you might use in your research and SSAP 24 problems.

**Dr Clacher:** Yes. But it is also about striking the balance on the portfolio investments as well.

**Mr Ternent:** I think I'll sum up what I am saying, and this includes the point about where the insurance company actuary is sitting with Solvency II. Each of these positions represents an insurance problem. What level of reserves are you prepared to aim to provide? The intermediate position – running the scheme as an ongoing fund – translates the persistence involved with running the scheme for 60 years (until everybody dies) into it having some kind of intrinsic investment value of its own.

**Dr Clacher:** That is good for the shareholders and the company, but it is problematic for employees and trustees.

**Mr Connolly:** Can I ask what the next step for this research is? Is there one?

**Dr Clacher:** We were thinking about looking at copulas in a separate paper. There is much debate around copulas and their application in pensions. Just now, we have to write a report for the profession. We also have to write a report for the Rotman International Centre for Pensions Management and submit a paper to the European Journal of Finance by September.

**Mr Freeman:** We also need to work on the liabilities. This has been our central problem: do you have to have stylised liabilities?

I think for the academic audience there is a case for keeping it quite stylistic. If we are going to do anything of any practical relevance we have to incorporate more complex liabilities.

**Mr Slee, FIA:** This is a frequentist approach using data from 1970 to 2010, and the approach included 2007 and 2008. If you had done it from 1966 to 2006, say, and you were doing this paper at the end of 2006, what would it have said about what was going to happen in the next two years?

**Mr Freeman:** The numbers would have been different. It would have still said that the funding risk is higher than most comparable models, but it would not have reached a level higher than the actual outcome.

**Mr Slee:** So if you had been using this model then, before the crash happened, would people say “What was the point of that model? It has not shown me where I need to be.”

**Mr Freeman:** But it would still have been a warning model. Our advice would have been “You are 20% solvent now but there is a good chance that in six months you will not be, or in 24 months you will not be.”

The probabilities will have been assessed at 10% to 15% whereas, with the benefit of hindsight, it happened with probability one. So we would not have got the risks high enough but we would have got the risks better than most of the alternatives available in 2008. As has been usual in my career, I would have been the bear.

**Mr Connolly:** What this paper shows me is that you could argue that it is futile to try and model extreme outcomes because the likelihood depends on the model you use.

There is an argument to say that it is not so much how often these events occur, but what you can do to protect yourself against these events. Have you any thoughts on that sort of logic?

**Mr Freeman:** I was talking to a risk manager a few weeks ago. We were having exactly this debate. Why model tails when tails are so difficult to estimate? His argument was that some metrics are more stable than others. He was a great advocate of value-at-risk rather than expected shortfall, as it is a much more stable measure. Expected shortfall often depends on what happens in the very end of the tail, which is very difficult to estimate accurately.

This is a value-at-risk exercise. So is it model-dependent? Yes it is. But what we have not modelled here is the expected shortfall.

**Mr Connolly:** Completely separate to the modelling, I would note the comment at the start of the abstract about the impact of quantitative easing on bond yields.

There was a speech by Charlie Bean a week or two ago from the Bank of England saying quantitative easing is not as bad as you might think. For example, it was not that bad for other asset classes, such as equities, where quantitative easing (QE) had kept values at a reasonable level. Have you any thoughts on his analysis?

**Dr Clacher:** Many! The problem we have is that there is a number of events which had an impact on yields, but QE clearly had an impact on the market. People then started seeing the UK as a safe haven because of the Eurozone crisis and so a lot of money has left the Eurozone and been invested in UK government securities. Therefore, quantitative easing has not had as significant an impact as it appeared. That is part of the problem. QE has clearly had an effect, but it has been compounded by other factors. More importantly, there has been nothing done to offset the impact that this is having on pensions and pension funds.

**Mr Freeman:** This is not my area of expertise. I would not feel competent to weigh up the different effects.

**Mr Slee:** Thank you for coming along tonight and giving an excellent presentation and thank you everybody for the ensuing debate. We look forward to seeing the final paper. If anybody subsequently decides that they had some points that they wished to make, please forward them to the Profession.

May I ask you all to thank Dr Clacher and Professor Freeman for their presentation and for the work they have done so far.