Mortality and Longevity – Making Financial Sense of the Highly Uncertain

-----

One Day Seminar

Wednesday 28 June 2006

Barbican Centre, London

Chairman's Welcome

THE CHAIRMAN (Mr Dave Grimshaw): Welcome, everybody, to the Actuarial Profession's first mortality and longevity seminar. This seminar is quite unusual in that it was sponsored and arranged through the education and CPD committees of both the Life and the Pensions Boards, which I think makes it highly unusual, if not unique, in recent years. It clearly has elicited a huge response - it is to be hoped not just because the CPD year ends this weekend! That resulted in a change in venue. I assume that the message got out clearly to everybody who is here. If there is anybody wandering around Staple Inn, obviously it was not quite so successful!

One of the hopes that was expressed when the seminar was first mooted was that we would achieve a mix of life and pensions actuaries. Glancing down the Delegate List that has certainly been achieved. Who knows, you may even be sat next to somebody you have not seen since you were in the exam room.

Mortality and longevity is of course a very topical subject, whether in the context of insurance companies pricing and reserving for annuities, a private sector pension scheme funding or of wider social issues, such as retirement ages in the public sector, and barely a day goes by without a news headline that somehow or other relates to mortality and longevity.

What we are hoping to do today is to cover a breadth of topics. I think the original programme of the seminar envisaged dividing it into three distinct segments. The first two sessions that we have this morning are on the theme of understanding recent experience. We have

Angus Macdonald and Brian Wilson talking to us before coffee. As with some of the other subsequent sessions, my intention is that we will try to run those two more or less together and then take questions at the end of those two presentations. We hope that we have allowed plenty of time for questions and discussion. Please do not go away thinking that there are unanswered questions. There is time to ask those.

After coffee, we move on to mortality projections, and before lunch we are looking very much at methodologies, a variety of methodologies, and very much at the statistical side of things.

After lunch Richard Willets is presenting from a different perspective looking at what intuitively one might expect to happen perhaps from looking at medical advances, etc.

Then, in the afternoon we move on to what do we do with all this information? How do we take account of the uncertainty? We have a pair of presentations focusing on communicating uncertainty in a pension scheme environment and in a life company environment.

Then, after afternoon tea, we have a view from Nigel Bankhead from the Board for Actuarial Standards. The last session of the day was a late change to the programme and perhaps fits better with the mortality projections, but we are going to have it at the end of the day. Steven Haberman and Richard Verrall, from Cass Business School, are going to present on modelling and measuring uncertainty with a focus on Lee-Carter projections.

It is a great pleasure to introduce Angus Macdonald. He is Professor of Actuarial Mathematics at Heriot Watt, having previously worked for Scottish Amicable. Angus qualified as an actuary with the Faculty in 1984. He is a member of the Faculty Council and a director of the Maxwell Institute for Mathematical Sciences at Edinburgh and Heriot Watt Universities. Angus is deputy chairman of the

CMI Executive Committee. He is speaking to us today as chairman of the CMI's Mortality Committee.

Understanding Recent Experience:

CMI Mortality Experience

PROFESSOR ANGUS MACDONALD: I should like to talk about some of the recent CMI experience, but also some of the work that the CMI has been doing recently, because it has been spending a considerable amount of time looking at questions of projections, which is also going to be discussed in other parts of this meeting. For that reason I am not proposing to say anything about certain of the CMI's experiences, like the assured lives, but to concentrate on longevity, because I guess that is where most of the interest in this meeting is centred.

I will say a little bit about the recent CMI experience and the new mortality tables that we are just about to publish; and then a little bit about the mortality improvements, which of course leads into the question of mortality projections; and just say a very brief word about what might happen next.

The basis of our current work is the changes which have taken place in mortality over the last ten years or so, indeed in the very short period of time since the 92 series of tables was published. I will illustrate this with the Life Office Pensioners' experience, comparing the most recent experience that underlies the forthcoming 00 tables with the 92 series, the previous standard tables, as they were projected at the time.

What this shows for males is the actual over expected mortality from 1992 to 2003, the actual being the experiences since 1992 and the expected being the standard tables published then. There are very clear indications, as is well-known, that the experience has out-run the projections that were made at that time.

For females, again showing both amounts and lives, the projections seem to have been somewhat closer to the outturn. But, nevertheless, that still indicates a fairly rapid rate of improvement since that was incorporated in the projections.

So, on the basis of the experiences since 1992 we have worked on a new set of mortality tables. You can see on the screen the complete list of tables. There are several new features of these tables that have not been in previous ones. In particular, the assured lives and term assurance tables are now in a non-smoker and smoker basis as well as on a combined basis, and we have personal pensioners and retirement annuities which are now taking over from the venerable immediate annuitants' experience. In fact, the immediate annuitants' experience showed some rather odd features that suggested the market there is changing significantly. We think that experience will no longer be graduated in future.

The work that has been done on the proposed tables had been exposed in the form of two working papers, working paper 12 on assured lives, and working paper 16 on pensioners and annuitants. We received feed back on these tables and have incorporated such feedback as we could. The final tables are due out shortly. I will say a little bit about that later on. We are, we hope, approaching the final approval process for these tables. FIMC has the role of approving the tables.

This is a very important change from the past: we used to obtain approval from the profession and then publish what we called standard tables. We now feel that it is not appropriate to do that since the job of setting standards has been passed over to the Board of Actuarial Standards who may or may not adopt any CMI tables for their own use. What we are now asking the profession to do is to approve what we will call base tables. We will scrupulously avoid in future the terminology 'standard tables', and leave it up to the BAS.

The comparison of the pensioners' tables with the 92 series, the last lot of tables projected up to the calendar year 2000, is quite interesting. Male and female lives are also in here. What you see on the right of the screen are the older ages. They are not too far off 100%. The big change is at the younger ages. That results from a decision we made this time. It has been a feature for many years, even for the experience which nominally represents the retirements at normal pension age, that the experience below age 60 tends to rise sharply. Clearly, normal retirements, as so labelled, are not all normal retirements.

In the past that experience has been graduated out. This time we left it in in response to what we thought were the needs of some of the users of the tables. So these spectacular apparent increases in mortality below age 65 are in fact simply the result of leaving that feature in the graduations this time that previously had been artificially removed.

The big question surrounding this experience is what mortality improvements are emerging from it. If you look back as far as the 1920s, which is when the CMI first began to collect statistics on a continuous basis, we can see remarkable improvements but varying greatly with age. Up to the 1950s there were great improvements at the younger ages and not very great improvements at the older ages.

Since 1950 the improvements at the younger ages have levelled out but the improvements at the older ages have now started to kick in. What you can see from the vertical dotted lines in the diagram are the quadrennia during which the data collected went to form the standard tables, A67, 70, PNA 90, and so on. So those give a snapshot of the CMI experiences that were used to form the standard tables and how the improvements sit alongside those periods.

Richard Willets may say more about this in the afternoon. If you look, in particular, at the infectious diseases as a cause of death then it is quite obvious where a large part of that improvement has come from between 1901 to 1910, and 2001. At young and young adult ages the enormous proportion of deaths due to infectious diseases has dropped away now to a very small proportion indeed.

These figures you can see on the screen are taken from the paper of Mr Willets.

At the older ages what we are seeing is where the scope for improvement possibly lies is in the current major causes of death at these ages: cancer, heart disease and respiratory disease. There is considerable discussion on how much of the improvement is explained by changes in smoking patterns. There have been changes in smoking patterns which may have caused a reduction in mortality after some period after smoking, and in due course -- because they take time to work through -- that may have, or have had, an effect on mortality due to lung cancer. But there is a great deal of discussion and research going on on the causes of mortality. That is a topic that I think Mr Willets is going to deal with this afternoon.

So the question is whether these improvements will continue or whether they will cease. The interesting feature is that the demographers cannot agree among themselves on this point. There is one group of demographers, and the most prominent is perhaps Jay Olshansky, which argues that these rates of improvement cannot continue, and they will level out and we will see a plateauing of the curve of deaths, and the projections should not continue to assume the current levels of improvement indefinitely into the future. Another group, which includes James [?] , and Shripad Tuljapurkar, believes that lifespans can increase indefinitely. They claim in defence of that view that there are drugs and therapies as yet undreamed of that will be based on fundamental discoveries about the mechanisms underlying ages.

I do not know how many of you have seen or heard of the Cambridge gerontologist, who often appears on radio and television and has indeed appeared at a CMI meeting in the past, whose prediction is that the first person who will live for 1000 years has already been born.

So there is no consensus about whether these rates of increase will continue or not.

Coming onto the mechanics of mortality projections, what was introduced with the CMI standard tables as far back as 1980, were the two-dimensional tables of mortality rates by age and calendar year, if they are arrayed as they are in the slide which is on the screen, in which one can pick out the individual cohorts of birth years as diagonals in this table, as indicated by the red dots. We have now become quite accustomed to the two-dimensional tables and the projections based on them.

From the point of view of the work that is going on just now, the data that we have to deal with are mortality rates themselves, q(x), rates of mortality or possibly forces of mortality, which tends to be the basis of the models these days. We have assured lives mortality, which is the protection and savings products from 1947 to 2003; we have UK population mortality from 1960 to 2003 by individual calendar year. What we do not have, unfortunately, because none of the experience is mature enough, is the annuitants' and pensioners' experiences, which is what we actually would like to project. So almost inevitably, we are going to base quantitative projections on the large and old data sets that we have, and then in some way apply the results to the pensioners' and annuitants' mortality that we need.

Before we talk about projections, the first thing that the CMI did was to try to pick out what patterns there might be in that past experience. That involved smoothing. The method that was chosen to smooth these data was two dimensional P-splines. I am not going to say anything technical about P-splines, but the author of this work, Dr Iain Currie, is here and is speaking later on, so if anybody did have any technical questions, now would be a good time to grab him.

We tried to pick out the patterns in the mortality rates. If you look at the mortality rates themselves, you can see some gross patterns like the patterns with age and general trends with calendar time. They are rather difficult to pick out. If instead you look at the improvement factors -- it is improvement factors of this kind which underlie the projection methodologies which the CMI had used in the past -- then some more interesting features emerged. This is the origin of what has now become known as the cohort effect.

If one picks out the rates at which mortality improvements are changing, where they are highest and where they are lowest, then what we find is that mortality rates are improving fastest, it appears, up the diagonals of these diagrams where we have calendar year along the bottom and age up the vertical axis. Birth cohorts in these diagrams are represented by diagonals.

You can see the pink and the purple lines representing the most rapid rates of mortality improvement. Assured lives, males and UK population males. Obviously, with much larger amounts of data there is a greater amount of detail that can be picked out for the UK population. But in both cases there is a consistent pattern of improvements by cohort. Of course, that led to the CMI producing the short, medium and long cohort projections back in 2002.

That kicked off a lot of work on projections in the CMI, starting with working paper 1, which introduced the cohort projections. Working paper 3 was intended to start off a discussion of different methodologies for projecting mortality. It was not a research paper. It was meant to inform the profession and the users of mortality projections because it was clear that the new feature that was coming along that was introduced in a very ad hoc way by the cohort projections, was uncertainty around mortality projections. Whatever mortality projections you choose to use, you can be sure that there will be a great deal of uncertainty about them. Just consider the views of these two opposing camps of demographers. Our concern was whether it was possible in any way to quantify the uncertainty around a mortality projection in a better way than we had done with the cohort projections in working paper 1.

So we received feedback on that working paper and then produced another one in which we went into a little bit more detail about two particular methods, the P-spline method and the Lee-Carter method. In doing so, we tried to make it very clear -- and I should like again to try to make it very clear -- that the CMI does not see itself as recommending any particular basis for mortality projections. There are several research programmes going on into mortality projection along different lines, many of which would be very promising or interesting or useful, depending on the purpose you had in mind. It so happens that these are the two methods to which the CMI devoted its limited time. But we in no way are suggesting that these are better than other methods or should be adopted as standards.

We issued in March 2006 working paper 20 which was a more in-depth exploration of the P-spline model. That has led to some statements that the CMI is promoting the P-spline model. I must tell you that that is not the case. It is simply the fact that the research has taken longer to complete on the Lee-Carter model, and in due course, when that is complete, we will publish another paper. We are not recommending the P-spline model.

I mentioned that we are seeking the Profession's approval of the basic graduated tables, and we will now call them base tables instead of standard tables. I should say that we are not asking the Profession to approve any method of mortality projection. In the current climate that is a very active research topic and we do not think that would be appropriate -- it would send exactly the message that we do not want to send.

Since the work we have done to date has been on P-splines, I will say a little bit about what we have done, and the things we have tried to look at. Just think of P-splines as one particular methodology. We will be looking at the same questions when we publish the working paper on the Lee-Carter method.

There are several issues to consider. First of all, what do the projections actually represent, particularly when you are considering the uncertainty rather than the base or central projection? Since cohort effects are now recognised as possibly significant, there is a question of whether you should project cohorts or project on the old age-period basis. It is a question of males and females, which is mainly one of data availability, and similarly assured lives versus the Office of National Statistics, the UK statistics. There is the impact of choosing the precise subset of the data on which you base your projections -- for example, what age range do you choose? Do you go down to age 60? Do you go down to age 50? Do you go down to age 40, or whatever?

The question I want to concentrate on -- in many ways the most interesting one -- is if this method had been available back in 1992, and it had been used then, how well would it have worked moving forward up to 2004 when the experience is now known?

In view of the time, and to leave some time for discussion, I will go straight to the question of back testing back to 1993. I should like to spend a little bit of time on that. The other topics have received some exposure.

What the slide on the screen shows are projections from 1993 to 2011, I think, using the assured lives experience up to 1993 as a basis for fitting the model and then making projections from 1993 into the future using the P-spline model if it had been available at that time. The solid blue line is the central projection which, if you like, is the mean or expected value of the distributions that are being projected into the future, and the dotted lines are confidence intervals. They give a 95% confidence interval which, as you can see, is expanding gently around the central projection.

The quantities that are being projected here are not mortality rates but annuity values since the purpose was to see what the effect would be on the end product, if you like, of the whole process, which is annuity values. These are for males aged 60. I will not go into the details of the basis. Perhaps the most interesting point is along the bottom, the light blue line is the annuity value for each future calendar year on the 92 basis projected forward with the projection methodology and the two-dimensional mortality table that was produced as a standard table in 1992. What you can see is a significant difference in annuity values and a gently expanding confidence interval.

We can then look initially just for the mean projection. We can compare that with the actual experience that has emerged since then, which is called the wiggly 50% line. What you can see is the projection has stayed fairly close to the central projection produced by the P-spline model. The dotted yellow line is what would be projected from 2003, which was the latest data to which this was fitted.

In terms of the central projection, the P-spline model has been a relatively good predictor of future annuity values.

Since each time the model is updated, as each calendar year's data comes in in 1994, 1995, 1996, and so on, you can improve the fit to the past data because you have an extra year's data. Thus you can advance the confidence intervals as well to allow for the additional data that is captured each year. What you can see there is, not surprisingly, they march along with the central projection and they do not change very much. All that is saying is that you have a smaller confidence interval in 2004 if you have data up to 2004 or 2003 than you have if you have to project forward from 1993 to that date. The confidence intervals will then naturally begin to spread out again after 2003.

So what then happens if we superimpose the cohort projections that we produced in 2002 on these Lee-Carter projections? We had the short cohort, medium cohort and long cohort projections, each individually on a deterministic basis. These were not stochastic at all. The short cohort obviously gives the lowest annuity values; it predicted the cohort effect to peter out soonest.

At the bottom of the slide the purple line shows the short cohort. The black line shows the medium cohort and the green line shows the long cohort. It is not too far out of kilter with what we see in 2000 going on to 2003-2004, but, looking beyond that, the cohort projections appear to be significantly less steep than the trend suggested by the P-spline projections. That is for males aged 60. We can look at similar effects if we look at higher ages. You can see on the slide the projected values for males aged 65 -- that is whole of life annuities sold to males aged 65; males aged 70; and males aged 75.

There appears to be somewhat better congruence of the cohort projections and the P-spline projections looking into the further future at the higher ages; somewhat less congruence at the younger age, aged 60.

Looking at some numbers of annuity values, there is quite a lot on this slide so I will try to explain what this is showing us. The figures to concentrate on are the percentages which you can see on the slide in red. From left to right we have males aged 60, 65 and 75. We are looking again at annuity values, whole of life annuities sold to those males. The grey boxes show what might be done on the current basis with a PMA 92 projected for annuities sold in the calendar year 2004 with the medium cohort projection.

Below that we have comparisons with two bases. First, with the age cohort projections; and, secondly, with the age period projections. The three lines show the 50 percentile projection and the two projections that form the confidence intervals, the 97.5 and the 2.5. "ac" stands for age cohort. "ap" stands for age period. These are comparing the two different ways in which the P-spline model can be used to project mortality rates forward.

The percentages compare annuity values on the respective P-spline bases or their confidence intervals with those in the grey boxes, what might be used under a current basis with the medium cohort projections.

What we can see is the differences are not so very large. The differences are not exceeding about 104% and the confidence intervals, with one exception, include 100% -- in other words, no difference.

Instead of seeing what impact this has on annuities sold in 2004, if we project even further forward on both bases to see what would be the impact on annuities assumed to be sold in 2030, we can see the impact of the P-spline model and the medium cohort model gradually drifting apart. 2030, of course, is a very long period to project forward into the future. That is the point of sale. You then have projections for the entire lifetime of the annuitants. These are extreme examples to show. On that projected basis annuities sold in 2030, the differences between P-spline methodology and the current methodology incorporating the medium cohort are much more significant.

From the standpoint of the present, that would matter for certain classes of business, but not others. But it does indicate that the medium cohort projections are not necessarily conservative compared even with the central P-spline projections or the most favourable part of the confidence interval of the P-spline projections.

That is a brief look at some of the work that the CMI has been doing recently. We hope, as soon as the work can be done, to be able to compare this with figures produced using a Lee-Carter model.

Uncertainty is really what is driving all this work, and also, I guess, what lies behind this seminar. I should like to say a word about the sources of uncertainty, what we think we are looking at and what we know we are not looking at.

What I have not shown you today, except in the last two slides, perhaps, if you dignify the cohort projections with the name of model, is we have not really looked at model uncertainty, comparing the P-spline projections with the Lee-Carter projections or projections produced with any of the other models that you will hear about today. That would give us some idea perhaps about the impact of model uncertainty.

We have addressed parameter uncertainty, since that is the entire basis of the projection model based on P-splines, and that is also what we are addressing in our work on the Lee-Carter model. In fact, that is what is taking up most of the time. We have not looked at stochastic uncertainty, which is simply whether the individual to whom you sell an annuity happens to live or to die, but that could easily be incorporated on top of the projections that we have here. In addition, there are sources of uncertainty that we could only measure if we collected more detailed data in order to be able to measure them.

For example, measurement error in the collection of the basic mortality statistics; unknown heterogeneity in the data, of which I would pick out in particular the effect of socio-economic class or affluence which we get only a very distant view of in the CMI's experiences to date because of the way the data are collected. We hope to do much better in future because we are intending to move on to a per-policy basis of data collection.

Then there is the whole underlying assumption that the past is going to be a good guide to the future. There are many reasons why that could change. Change in business mix is one example, and you can think of many more catastrophic examples affecting society that might invalidate that assumption.

So, where next? Our next steps, first of all, are to get the graduated tables into the public domain with the approval of the profession, not standard tables any more but base tables, as I have said. We are hoping that FIMC will be able to complete the approval process in July, if all goes well.

On the projections side, the work is continuing on the Lee-Carter model, and as soon as we have accumulated enough to publish on that we will produce another working paper along the lines of the one that we produced on the P-spline. Apart from that, the CMI will obviously continue its basic work of data collection and data analysis and will continue to be an interested participant in the research process which is currently underway into the whole topic of mortality projection. I shall be very interested to hear about that other work that is underway in other methods that are going to be discussed later in this meeting.

Thank you.

THE CHAIRMAN: I will now introduce Brian Wilson. Brian is a research actuary at Hewitt Associates and his work includes the production of technical support material for both consultants and clients. He has been with the firm since 1976 and specialises in advising corporate clients on all aspects of pension schemes. Brian is a member of the Legislation and Technical Support and Research Committees for the Pensions Board of the Actuarial Profession. Perhaps more pertinent to this morning's session, he chairs the Working Party on the Mortality of Pensioners in Self-Administered Pension Schemes.

Understanding Recent Experience:

Self-Administered Pension Scheme Investigation

MR BRIAN WILSON: The SAPS Working Party's work is relatively new. In fact there have been one or two developments in the past couple of weeks which means that one or two of the slides that I am going to show you may be out of date. I will fill you in when we come to it.

The history of the work is that the profession looked at a pilot study covering the years 1998-2000 in order to ascertain whether the life office pensioner data was valid for pension schemes and, if not, were there good grounds for starting up a separate study of self-administered pension scheme pensioner mortality? That did indeed indicate that there were differences and that something else should be started.

In 2002 the CMI was commissioned by the Actuarial Profession to begin its investigation. The reporting line, however, for the work was via the Technical Support and

Research Committee of the Pensions Board. The CMI were doing a lot of the legwork but we were not reporting to the CMI.

The working party that was established has four people on it: myself, Nigel Bodie, Jonathan Lawlor and Andrew Gaches. We will be hearing a little bit more from Nigel later on this afternoon.

The slide you can see on the screen is now out of date because a fortnight ago the CMI Executive Committee formally established the SAPS Working Party as a Committee of the CMI so we now report jointly to the Pensions Board and to the CMI Executive Committee.

How large is the study? So far data have been submitted from 11 consultancies covering 255 schemes with 2.67 million records. We have been focusing on the years 2000-2003. We have produced three working papers, Nos 4, 9 and 17, and which can be downloaded from the Institute website.

To look at those in slightly more detail, the first one which we produced in March 2004 analysed the data collected between the beginning date of the study, which was 2003, and March 2004. So there was 15 months' worth of data collected there. Clearly, the data that we were collecting was largely from triennial valuations so it covered the three years up to the date that the data was collected. In working paper No 9 we analysed that in more detail.

Working paper No 17 looked at the data collected up to May 2005. So we have almost 2½ years' worth of data collected there. The working party has been intending to produce a more detailed analysis of that data to publish later this year. However, we are now questioning whether, with another year behind us, we might actually look at the data collected up to May 2006 and compare that with the working paper 9 analysis. We will have twice as much data and we will look at female rates, whereas we were concentrating more on the male rates in working paper No 9.

So, what did we analyse in working paper No 9? We looked basically at three things. We looked at amounts and we looked at four different amounts bands. We had some difficulty deciding how to split the data. Initially if you split it into four lots of 25% by amounts going up, we discovered that the bottom two were very, very similar indeed. So we decided then to look at the bottom 50% together, the next 25%, and then the top two were the top two 12½% bands. That is, with a little bit of rounding, how we looked at those.

Females are going to be a lot more of a challenge because there is not the spread of pension amounts for females that there is for males. But we are going to look and see what we can discover in that.

Data are split by industry classification. Thirteen different classifications. There was sufficient data in working paper 9 to give significant results with seven out of those thirteen. Later this year we hope to be able to look at rather more. We compared ill-health pensioners against other pensioners.

The slide you can now see is quite interesting. Again, there is a slight confusion here because CMI on this refers to the CMI insured pensioner investigation that Angus has been talking about. SAPS is the self-administered pension scheme investigation which now is of course a CMI investigation as well. So CMI equals insured and SAPS equals self-administered.

This slide compares the volume of data. We are looking at deaths in order to compare that. We could have looked at exposed to risk. You will see from the pale blue lines on the left of the slide that by working paper 17 up to May 2005 we had, over a four-year period, which compares to the CMI's four-year period for the LO table data, a third more males by lives.

If you look by amounts, the next two dark blue lines, you will see that we have significantly more by amounts, so if you want a pension come to a self-administered pension scheme! Then that tall blue line in the middle of the slide is broken down by the next three lines. The data is submitted to us by schemes. Some schemes can separate out normal retirees from ill-health retirees and others cannot. The one that says “combined” is from schemes that do not have the data to split them out. The other two are from schemes that do have the data to split them out which enables us to do various comparisons, just to look at the ill-health experience compared to the normal.

I do not have a slide which looks at the females, but that looks even more heavily biased in favour of the SAPS investigation. I have some numbers in front of me. On a lives basis, the insured investigation has 18,000 deaths. The SAPS has 73,400 deaths, which is about three-quarters the size of the male investigation. But it is four times larger than the insured pensioner investigation.

If we look on an amounts basis, then the CMI pensioner 00 investigation has £15 million per annum deaths, the

SAPS has got £166 million, which is only a third as many as the men but is actually 11 times larger than the CMI insured pensioner data. So we are already a very large investigation in CMI terms.

How do the results look? This looks at normal retirements on a lives basis. We are comparing against the red line the PML 92 short cohort projected up to the year 2000, and then the blue line looks at the draft 00 table numbers that have been published. You will see that, apart from the very left-hand, the rates are coming out as heavier than the insured pensioner rates. We are above the 100% actual over expected.

The divergence at the bottom of the slide is because of course the 92 series and the 00 series themselves diverge, as was explained by Angus. We ask: what is the reason for the heavier mortality? Of course, we cannot be absolutely sure but one of the criteria for submitting data is that there are a minimum of 500 pensioner lives in the data that we collect. If we were to collect data from all schemes, the CMI staff would be totally inundated and the work would be impossible to do. So we had to have reasonable cut-off to keep costs within bounds but it means that we have a larger number of the old large heavy industry type schemes in there than the average if we took the whole totality of the self-administered scheme population.

That is one reason why we split the amount into four bands so that we could get a better handle on exactly how the mortality did actually compare.

The slide that you now can see is exactly the same as the last one except on an amounts basis.

This next slide is for females on an amounts basis. You can see that there is a consistency here that they are all just running over the 100% actual over expected.

This is where it gets more interesting because on this slide you can see the lines which show the four different amount bands that we have divided the data into. The bottom band is the red line up to £4,500, which is the bottom 50% of data. So the average comes up at £4,500. The blue line is from there up to £8,500. The third band is £8,500 to £13,000 and the top band is above £13,000.

You can see that there is a clear divergence of the four bands. The bands are compared against the 92 series with a short cohort projection, and the top two bands are below 100% so have lighter mortality, whereas the bottom two bands are above, the blue one very, very close. It is just marginally above.

Although we are at a very early stage in the investigation, we asked ourselves whether we had actually noticed any difference on a year by year basis. The slide you can now see is looking at all male lives and the actual over expected deaths of those four amount bands in the years 2000 to 2003. There are some interesting features there. You can see that they consistently remain separate from each other, as one would expect. They appear to be converging up to 2002 and then staying steady up to 2003. We have no idea what the reason for this is. If anybody can tell us, we would be interested to know. It is probably just too early in the study.

The next slide shows a comparison. I have referred to "graduated". Interestingly, I took the programme for the P-spline projections. I fed the raw data in and projected it four 0 years which very conveniently gives me a graduation of the data. Not the best way of doing it, probably. But, nevertheless, it shows, looking at the male data against the draft 00 series, you can see consistently heavier throughout, apart from the very, very young ages, where we do not have much data anyway.

Again, this slide shows the four different amount bands compared to the all data band. The all data one is the darker blue line, which is the second one from the top of the graph. As you can see, the separation from the ‘q’s at the amounts occurs at all ages up. Funny things happen round about age 63, which probably is to do with the graduation rather than anything else.

Nigel Bodie will be interested in this next slide, because he produced it. It looks at the logs of qx over px which theory tells us should progress in a fairly linear way. We have some quite nice straight lines there.

When we look at it on the next slide, on amounts, again and we come up with some quite nice straight lines. We will be looking at graduating the data, what I might call “properly”, using the latest CMI software as has been used for the 00 series. Once we have had a good look at that, we hope to publish some graduated rates by the end of the year.

The next slide shows how I have imposed the medium cohort projection on top of the data as graduated roughly by me just in order to get some idea of the life expectancy of the four different amount bands. You will see that the difference between being on a low pension and on a high pension can add 3½ years to your life.

In terms of where we are headed, I shall explain the diagram that you can see on the screen. The year data was collected is on the left-hand side. We started in January 2003. In that year the data collected related very largely to exposed to risks in calendar years 2000, 2001 and 2002. You can see the horizontal red line. The following year we move everything along one to the right, and similarly for 2005 along to the right again.

You will see that because we have largely triennial valuation data, the first year for which we will have a complete set of data is 2002. The second year for which we will have a complete set of data is 2003 and then 2004 and 2005. We will not get that 2005 data set complete until the end of 2008. If we were to follow the normal CMI pattern of using four complete years' worth of data in producing base tables, then we would not be producing a base table until 2009.

However, we will have additional data from the earlier years and additional data from the years beyond 2005 at an earlier date than 2009, and we might consider looking at producing some base tables next year. That is all subject to what the data actually looks like when we come to look at it. If we think that there is anything “funny” about it then we might need to wait a year or two more. However, we live in hope that we will be able to do something sooner rather than later.

THE CHAIRMAN: Thank you to both speakers for their presentations. May I ask for questions now?

Questions

Mr [?] (Hannover Life Re, Dublin): I am curious. Would you guys have more information available attached with the SAPS experience that might allow you to carve it up by different factors, so not just by sum assureds? Would you be able to carve it up by post code? What else do you guys have available? It would be interesting to know.

MR WILSON: We collect industry classification so that we can carve it up by industry. I indicated that we had had a first stab at that. Unfortunately, we do not have post code data. The data sent in for actuarial valuations by companies does not normally include the address data of pensioners. We might start asking for that so as the study progresses we will be able to get something rather better on the postcoding.

The other thing that we might look at is geographical area where again we are not collecting data on that at present. We might also look in a rather more sophisticated way at amounts because the smaller pensions can be coming from people on quite high salaries who change jobs frequently, so they build up lots of small pensions over their working life. So getting a pensionable salary at retirement as an additional data item would also be very useful, and we might consider doing that in future years.

These are lots of interesting things that we might be doing. Unfortunately, the more data that we are asked to collect, the less one gets with a complete data set in it and the more cost there is in actually processing it. But we live in high hopes of being able to do better things in the future.

PROF MACDONALD: Post code is one of the items that we are hoping to collect when we switch to a per policy basis of data collection. The issue at the moment, which I think is almost resolved, is the amount of detail that we would be allowed to collect and not conflict with the approach to privacy of the Data Protection Act.

MR MATTHEW EDWARDS (Watson Wyatt): I have a question for Angus Macdonald. On your graph, the first of the graphs showing the P-spline backfitting against the previous 92 series, at the start the projection there was already quite a big divergence between the P-spline backfitting and the 92 series. I just wonder whether there is a particular reason why they already start off quite separate, even before they then begin to diverge further.

PROF MACDONALD: It is because the annuity rates in 1993 are not annuity rates based on historical data entirely. They are annuity rates based on the fit to the historical data projected into the future. So what that is showing is divergence between the P-spline projection and the 92 projection right from 1992.

I have not actually seen this, but I would expect that if you ignore the projections and just look at the calendar year 1992 fitted both by the CMI graduations and the P-spline model, I would expect those to be similar.

MR WILSON: I think that he difference, Matthew, is the original 92 projection took no account of the cohort feature. That only came in with the interim cohort projection so you immediately have a difference there in terms of what you are projecting forward, even though you are working from the same data.

MR IAN MORAN (Deloitte): I have a question for Mr Wilson. Can you just expand on the way that you selected the schemes that you based your exposure on? Obviously, it is quite a critical factor to explain what the results are.

MR WILSON: Yes. All the consultancies were invited to submit data. In order to contain costs, we had to limit the number of schemes from which we were collecting data. We selected those with 500 pension lives or more. Essentially, we have taken whatever the consultancy have sent that fit that criteria.

There are issues in that some schemes do not have data by amounts for pensioner deaths, and there we have not rejected that data but we have used that data and have put in various assumptions relating to amounts at death based on other information contained in the data.

THE CHAIRMAN: Thank you very much for your questions and thank you again to Brian and Angus for their presentations this morning.

Mortality Projection Methodologies:

P-spline / Lee Carter

THE CHAIRMAN: I should like to introduce Iain Currie, who is going to look at both P-spline and Lee-Carter. Dr Currie is a Reader in Statistics in the Department of Actuarial Mathematics and Statistics at Heriot-Watt University, Edinburgh. His main research interest is in multi-dimensional smoothing methods with particular applications to the modelling and forecasting of mortality rates.

His recent Royal Statistical Society paper "Generalised Linear Array Models with Applications to Multi-dimensional Smoothing" with Paul Eilers of Leiden[?] University, and Maria Durban of Carlos III University, summarises his work over the last few years.

So, Iain, I will ask you to take us through P-spline and Lee-Carter, please.

DR IAIN CURRIE: As an outsider, I feel that I should first of all thank Dave for the invitation to address you, and also perhaps clarify my relationship between myself, on the one hand, as an academic statistician, and the CMI and the actuarial profession on the other. I guess that the relationship is pretty clear cut and quite separate. As an academic, my work is available in the public domain, and the CMI, and indeed any of you, are quite at liberty to read it, like it or dislike it and make use of it as you see fit.

Today the purpose is to talk about one of the methods that can be used to address the forecasting problem.

I thought I would start with one or two more general remarks. Future mortality is obviously fundamental to much of actuarial work. The CMI, even as I speak, I suppose, are busy forecasting mortality, and I believe even up to 100 years ahead.

This is not an exact science. How can it be? Go back a 100 years, we are in the middle of the Edwardian age, the height of British imperial power, the First and Second World Wars, the rise and fall of communism, the age of scientific medicine. All these were in the future and all had huge effects on human mortality.

So who really knows what the future holds? What will be the effect of the rise of religious fundamentalism, the emergence of China and India as world economic powers? And over all of this sits the spectre of global warming. Yet forecast we must. All we can do is to take the past as a guide to the future. This makes a very fundamental assumption about some stability in the future course of mortality, that the past can indeed be used as a guide to the future. How good a guide is not clear but it is all that we have got.

So it is with this background that I intend to discuss just one of the methods and one of the ways of forecasting future mortality.

This slide that you can see on the screen is an example. I have one data set. This is the CMI assured lives data set. I have selected the subset of ages 20 to 90 and all the years that I had available, that is 1947 up to 2002, and we think of this as living in a data matrix classified by age and by year. The problem is to forecast some future date, and I have rather arbitrarily chosen 2046.

The way that I am going to try to present this is I am going to try to describe to you what the P-spline method is. To do this I am going to take a much simpler problem and just look at P-splines as a forecasting problem in one dimension. This will set up a general method. It is like a plumber, I am going to get this tool, this forecasting method, in one dimension, and then I am to take this tool and apply it to the forecasting and mortality table, and we will see that we can apply this in a number of ways, of which the P-spline method that you have had referred to this morning is just one example.

That is the second part of the talk. Once we have set that up, I will be talking about the Lee-Carter model and I will talk a little about another model that we have not heard anything about, the Age-Period-Cohort model, and then the 2-d P-spline method which we have heard a lot about already.

My main area of research is in the general area of smoothing. P-splines is one method of smoothing. Smoothing, of course, has been around for a very, very long time. One character who knew all about smoothing was a very illustrious forebear of the yours, Gompertz, who round about 1830 made the really remarkable discovery that if you plotted the log of the mortality rates against age, then lo! and behold, over a great spread of human life you got more or less a straight line.

You can see on the screen the startling discovery of Gompertz. This shows, startling round about ages 30 to 40 all the way up to age 90 we have the log mortality rates lying on a straight line. There is not anything smoother than a straight line. That is very smooth and you can just fit this with ordinary regression. It is not quite ordinary regression, but let us just think of it as ordinary regression. That is a very elementary thing to do. That is something that has been around for over 170 years.

We can apply this method, exactly the same idea as Gompertz applied. He was applying it for a fixed year for all ages. I am going to apply it here for a fixed age. In other words, across the table in an opposite direction. I happen to have chosen age 70, which is an age that you might be particularly interested in. Then we plot the log mortality rates against year. You can see on the screen the data going from 1947 all way up to 2000. You plot it against year and you can stick a straight line regression through that. That is just a mathematical exercise.

As a statistician you might look at that, or as anybody with a little bit of common sense, and say that you are not very impressed by that and you do not think much of that.

But even a first-year student would know how to fix that. Instead of fitting our straight line regression, we could introduce a quadratic regression. That appears to do rather well. That is the first theme of the talk in fact on this very simple slide that you can see on the screen. It is model selection.

Here I have two competing models, the straight line model and the quadratic model. All modelling and forecasting is all about choices, choices of method, choices of model, and here is a very clear example of where one model is rather poor and another model might be rather better.

So then you could imagine an argument like this. This is a very compelling argument. The quadratic model clearly is a very good model -- and this is all in inverted commas, you understand! It describes mortality, at least for age 70, very well over the last 56 years. So let us make the great leap of faith and assume since the quadratic model describes things wonderfully well in the past, it will continue to hold in the future.

You can now see the slide on the screen. That is exactly the same function just continued. For those of you of a nervous disposition, you have perhaps just had a heart attack because these forecasts are really outrageous, most actuaries would like to think. We are forecasting -- I am not quite sure what the figure is -- mortality at age 70 in 2046 probably dropping to about the mortality of a person aged 50 or even younger. So this is forecasting an absolutely huge improvement in mortality.

But then somebody else might come along and say, "I am not very impressed with that. I do not think that is a very sensible forecast at all. I think it will be much better if we stood at the present-day and instead of continuing the quadratic, we will just draw a tangent at that point. We can all differentiate the quadratic. We can work out the equation in a straight line and there it is." So that is another forecast. So that is the second theme of the talk.

We have had to choose a model and now we have to choose a method of forecasting. You can see that these two different methods of forecasting give very, very different forecasts. If you were to ask me now which I thought was the better forecast, I would rummage in my bag and see whether I had brought my crystal ball because we are not really going to know until 2046. It is a subjective business. You may like the blue forecast for business reasons, but that is not a very good reason for choosing it.

Let us try to think about how we might go about smoothing in a slightly general way. I want you to try to think about what regression is. Cast your minds back to those heady undergraduates days when you were asleep during the statistics lectures.

A linear regression is simply taking a linear combination of what are called basis functions. The basis functions are the unit and the 45 degree line, unit x. We are looking to choose a linear combination of 1 and x. I put up the example of the Gompertz one. Or you could take a linear combination of 1,x and x^2 or indeed any polynomial basis.

Then the way that the regression method works is it just says that in terms of the function it is the log of the mortality is linear; the log of the mortality is quadratic; the log of the mortality is polynomial. If you use the data to select the values of a and the b by some fitting process, then you will have a fit, and on the right-hand side you get the fitted log mortalities, and then you could use that to do forecasting if you happen to think that that was a good idea.

Regression here is being thought of as taking linear combinations of basis functions. Here are three bases. You can use any functions at all. There is nothing sacred about polynomials. So we are just going to put in any old function.

Before that, I should say that the underlying technology that we have going on here is that of the generalised linear model, all part of the actuarial exam system nowadays, so the model that is underlying all of this is to say that the number of deaths follows a Poisson distribution, and the mean number of deaths is the central exposed to risk times the force of mortality, and then the Gompertz type laws are that the log of that mean number of deaths is going to be linear. That is the underlying technology.

So, instead of using polynomials as bases we can, being a mathematician, think that there is nothing special about polynomials, we will use any old functions. I have called these functions b 1 up to b p. These could indeed be any functions. I am going to choose particular ones, but it does not actually matter. It could be anything.

The main point that I making here is that the model values are given in exactly the same way. You are going to model the log of the force of mortality as a linear combination of the basis functions, and you going to use the data to choose the particular linear combination that does the job best of all.

So what are we going to choose for basis functions? I am going to use things called B-splines. This is where people start getting nervous -- for the second time!

You can see on the screen a picture of a B-spline. Really quite a friendly looking function. There is one huge difference between a polynomial, on the one hand, and a B-spline here. The B-spline is zero almost everywhere. These are local functions, they only take non-zero values over a finite range. This particular B-spline has non-zero values between the ages of about 23-24 and about 77. Otherwise, it looks like a nice bell shape curve.

It looks very smooth. If anybody is able to say it is not a smooth function, I should be absolutely astonished, but it is not completely smooth. There is actually a discontinuity in the third derivative at these blocks. The human eye, of course, is quite unable to detect that.

This is a cubic B-spline. It consists of four pieces. Each piece is a cubic polynomial and they are bolted together in a cunning way that so that they are continuous. The first and second derivatives are exactly the same. That is why it looks nice and smooth. That is one of the basis functions.

You can see on the screen the complete set. One very nice thing about this, which is again very different from polynomials, is that the basis covers the whole of the age range, of course, and the great thing about these bases is that it looks exactly the same wherever you are. If you are at age 80 and look up at the basis, it looks exactly the same as if you looked up at age 20. This is quite different from polynomials which look very different in different parts of the age range.

The main reason why you should not use polynomials for forecasting is the fact that they are global functions and so their properties outwith the data are really rather unpredictable. We are hoping to get round that by using these local basis functions.

So, we have this new basis, we have some data, so all we need to do, it is an ordinary regression, there is nothing fancy here, is stick it into our regression package. We are feeling quite bullish about this. But, oh dear, it has not worked at all well. You have 23 of these basis functions marching across the year and it is obvious what has gone wrong. We have too many of these things. The thing about these B-splines is that they give you some flexibility so you can mimic what the data is doing locally. We seem to have rather overcooked it here.

One strategy might be to try to cut down the B-splines and even maybe position them rather more carefully. That is certainly a way that you can go.

The P-spline method uses a rather different strategy, which is based on the following observation, which my colleague Paul Eilers and his friend Brian Marx made about ten years ago. They made following observation. If you actually plot the coefficients belonging to each of these B-splines, the coefficients themselves sort of track the data. It is the erratic nature of the coefficients that transmits through to the erratic nature of the fit. What they said was, "Well, if we could some how or other smooth out the coefficients" -- you are not allowed to smooth the data; the data are sacred, absolutely sacrosanct, but you may choose the coefficients -- "in some cunning way then maybe we could smooth out this fit". That is exactly what they did.

This is the master stroke. You can see on the screen a measure of the difference between adjacent coefficients. The first term in that quadratic expression measures the difference between the first, second and third coefficient, and so on, all the way right across the whole range of coefficients and we are able to rewrite that as a quadratic form in theta, where P 2 is a second order difference matrix. We are using quadratic differences here.

This is what is known as a roughness penalty. The reason why it is called a roughness penalty is because the further the thetas are apart, the bigger the value of P 2 will be, and the rougher the fitted function will be. The closer the thetas are together, the smoother the function will be.

Now we have a way of measuring roughness. How do we incorporate that into the estimation process? We have sitting underneath all of this a generalised linear model. That is very easy because we know how to do that. We just use maximum likelihood. So, no problems there, but we have something which does not seem to gel terribly well with maximum likelihood so we have to have a way of bringing maximum likelihood estimation together with measurement of roughness. This is done with the concept known as penalised likelihood.

You can think of this as a kind of balancing process. The L of theta is the likelihood function. That measures the fit. The closer your fitted function is to the data, the bigger the likelihood will be. So it measures fit. On the other hand, the second term is our roughness penalty and it measures smoothness.

So we have two things competing here. On one side we have the fit as measured by the likelihood; on the other hand, we have the roughness as measured by the penalty. Then we have to make a choice, how much importance do you give to one side as opposed to the other? So there is a balancing factor here which is called the smoothing parameter, and that attaches a weight, you can weight the likelihood, it makes no difference, and you just weight it to the roughness penalty. You can think of this as a kind of way of deciding. If you have a very big value of lander that means you are going to penalise roughness very heavily and you will get a very smooth function. If you have a very small value of Lander, roughness is not a big issue for you and the likelihood can take over and you get a very rough function or one that follows the data very closely, as we saw in previous slides. It is a balance of fit and smoothness.

This is perhaps quite a lot to take in here, so the following might help you a little bit to think that there are two extreme scenarios here. At one end of the scale, with certain parameterisations, if we take no roughness at all we can get interpolation. At the other end of the scale, if we put in an infinite amount of smoothing, which says “any amount of roughness is anathema to me”, then we get back our old friend linear regression.

So you can think of this method as providing a continuum of fits which range from interpolation at the one end to linear regression at the other end.

So how does it work? You can see the slide on the screen. With a suitable choice of the Lander, I am not going to discuss that discuss that here, we still have 23 B-splines in the basis but we have applied this roughness penalty, and lo! and behold, we get something that we would be reasonably happy with. This has picked out the trend in the data in a fairly sensitive fashion.

I have also plotted the coefficients themselves. The data are the black dots and the coefficients themselves lie on the fitted line.

This is also going to tell us how to do forecasting. We can use this as a trick for forecasting. All we need to do is to say the following. [Indicating on-screen]: If we are here and here, where is the next coefficient? I wanted to maintain the smoothness, so if it is a quadratic penalty that we are using, this will just give us a linear interpolation. So that is exactly what happens. You essentially forecast the coefficients and that gives us the forecast for the actual function.

The message here is that a quadratic penalty gives you a linear forecast. Of course, you can use different penalties and they will give you different forecasts. Again, we cannot escape from this choice. If you want to have a quadratic forecast, you simply use a cubic penalty. If you want to have a constant forecast, if you want to make it go horizontal, then you would use a first degree forecast. You can take linear combinations of these as a well, so everything is up for grabs. Linear forecasting is both simple and, in some ways, it is quite appealing. That is the one that I have hung my hat on for this particular talk, and I think that the CMI have hung their hat on that particular peg as well.

Now you could say to me “This is quite impressive but where have these 23 B-splines come from? If I were doing that I would maybe choose a different number. Does that give you a different answer?" One of the very nice properties of the P-spline method is that it has a certain almost invariance.

You can see on the screen that with 23 B-splines the difference between the knots in each B-spline is about 2.2 years and nine months. We can put the knots further apart, and we could have the difference between the knots 11 years. Then we can go to the whole process again. I have shown both fits here and both forecasts. You cannot actually see the first fit that I gave you because the second fit is plotted directly over the top of it. You get the same answer.

So in terms of the data area, you get the same answer, and indeed the actual forecasts come out to be almost the same.

One of the key things about the P-spline method is that it does not really matter how many knots you use provided you used too many. If you do not use too many knots the penalty has nothing to work on. Too many knots and the penalty will work to give you a desirable answer.

That is the end of the first part of the talk. We now have a method assembled for smoothing our particular data in one dimension and we are going apply that in a number of different ways.

The first method we are going to use – there has been some mention of the Lee-Carter method – you can see on the screen what all the fuss is about. You can see a formula that says that the log of the mortality at age I and year J is equal to alpha I equals an age effect, plus a time effect kappa J, which is modulated by an age dependent factor.

I will say that again; there is a wee bit to take in. The log of the force of mortality at age I and year J is equal to the age effect at age I plus the time effect at year J, and that time effect is modulated by an age dependent factor.

If you have managed to take that in, good for you, but I think that that is pretty tricky, so let us draw some pictures.

This is a generalised linear model so you can fit this without any difficulty at all – well, it is almost a generalised linear model. You can fit it, and you can see on the screen what you get.

I will take you through this. The first upper and left panel gives you a plot of the As. That is immediately recognisable as aggregate mortality by age. We can understand what the As stand for. It is overall mortality. So the Lee-Carter model says that the log of the force of mortality is equal to some overall mortality for age and then you are going to change it by something that is time dependent.

Now, this is what all the fuss is about. This is why there are nearly 150 people sitting in front of the speakers today. Kappa is the time dependent factor. That is the bottom left-hand corner of the slide. You can see that it is heading downwards at an alarmingly increasing rate. That is saying that mortality is generally improving with time. The modulating factors show you how that time effect is moderated at each age. That is a somewhat more erratic function, particularly at the younger ages where it goes zooming all over the place.

You can fit that model, and you can see on the screen the fitted mortality for age 70. I think that this is a most amazing slide.

Those of you who are sitting at the back of the hall perhaps may not be able to read the scales. If you are unable to read the scales, then you will be saying to yourself, "Goodness me, that function on the left at the bottom of the slide is exactly the same function as on the right bottom." And of course they are. That is what the Lee-Carter model says.

It says that at every age the mortality looks the same except that you have to relocate for age and rescale for age. It is simply a rescaling and a relocating of the time factor. This is even clearer if you plot four ages, so this is 50, 60, 70 and 80, and you can see that it seems to track the data pretty well. But you get exactly the same function.

My take on this would be to say that this is probably not a very good model. It seems highly unlikely that the force of mortality will follow the same functional form at every age, which is what to the Lee-Carter model says.

So, It does not get many marks as a model, but it may get a lot of marks as a forecasting method. One thing that it does do, if you go back a slide, is that it is very clever.

The clever thing about the Lee-Carter model is that it takes a table which has to be forecast and it reduces this forecasting problem to one dimension. It is a real trick, we have turned a two-dimensional forecasting problem into a one-dimensional forecasting problem. To do that we have had to use a rather poor model. It does not mean to say that it is not going to give us a good forecast because the forecasting should be quite easy because we have to do it in only one dimension.

You can see on the screen the forecasts that you get. The Lee-Carter method says that you should look at the kappas and use a time series method for forecasting. These are time series forecasts of the kappas. You then assume that the alphas and the betas are the same throughout the whole of the data region and the forecast region, and you can see on the screen the forecasts that result.

That is the original Lee-Carter model. Now I am going to be the plumber here. I go into my toolbox and I pick out my smoothing method and I simply apply this smoothing method to the Lee-Carter model. I take the As and I say that I am going to make that a smooth function. You can see my B-spline basis which says that the As are going to be smooth. It says the betas are going to be smooth and it says that the kappas are going to be smooth.

You can see on the screen what you get. You do not really see much difference on the As because the scale is so large. You do not see any wiggliness there. These are nice smooth functions. The beta has been extravagantly smoothed. All the wildness of it has gone. The kappa has also been smoothed.

I guess this explains my general feeling about smoothing methods. If you want to forecast, it would be a good idea to get the underlying signal, and if you have the underlying signal exposed, it might be easier to make the forecast. That is the thinking behind all this and you can see on the screen the forecast using the penalties. All of this is done using our little tool.

You can see on the screen the smooth Lee-Carter. These are the forecasts.

Now, the age-period-cohort model. This is an additive model that says you have an age effect or a period effect and a cohort effect. Again, this is a generalised linear model. It is very easy to fit. You notice that it does not have the problem that the Lee-Carter model has. It is sensitive to the extent that it will fit different functional forms at different ages. There is local recognition of what is going on in the mortality table rather than in this imposition of the single form over the entire table. That is the discrete age period cohort model. You have a parameter for every age, a parameter for every year and a parameter for every cohort. There are lots are lots of parameters here.

We can apply our plumber's technique again. You can take the alphas and smooth them; you can take the betas and smooth them. You can take the gammas, the cohort effects, and smooth them.

You can see on the screen that I have shown both the original age period cohort with its huge number of parameters and the smooth age period cohort with its smooth parameter set. Of course, the smooth functions track the discrete functions, as shown.

One of the things that you can do here is you can decompose this smooth model into different components. This slide shows a planar component. Imagine you have a mortality surface. You just take a plane out of that and then three components which describe the additional age variation, the additional year variation and the additional cohort variation. You can see summaries from that model. These are the improvements by year for different ages. These give you a summary of the improvement in mortality which is saying that the improvement in mortality is roughly .015 per annum. That comes out of the age period cohort model. You get the same kind of things for components of year. That is going from age 20 to age 90. This is averaged over the cohort effect, and finally the components of cohort.

This is still work in progress. I had hoped to finish this for this talk but we are still working on this. We have the model, the work, the smooth models in place and the next plan is to do the forecasting but that has not been finished in time.

Finally, I want to say a little bit about the 2-d process. We have our basic tool bag here. That says how you can smooth in one dimension. I know how to smooth an age, I just put a B-spline basis along the ages and smooth. I know how to smooth the years. I put a B-spline basis along the years and smooth. The question is: can we actually stick these things together? We call these the marginal bases, the basis for smoothing age and the basis for smoothing year.

We put these together to give a way of smoothing a table. Each of these B-spline bases will only smooth an age and a year respectively. Can we smooth the whole table? We have somehow to glue them together.

We can do that, and I will show you the slide of the gluing process. You can see on the screen what you get when you glue them together. I think that is clear. You can see the marginal basis. You can see the one-dimensional basis running one way and another one-dimensional basis running the other way and you run the two things together. That gives you a two-dimensional basis, so instead of a sort of cardboard cutout which you get with one dimension, you have two-dimensional smooth hillocks. These are two-dimensional B-splines. The method then is applied in exactly the same way.

The coefficients are attached to each of the B-splines and then you have to say to yourself how do you make these smooth? One way, obviously, of doing it is to smooth along each of the axes. We smooth along the years and we smooth along the ages. We hope by that way to get a smooth surface. You could smooth along different axes, I guess, but this is by far and away the simplest way to do it so that is what we do.

This is ordinary regression. Basically, there is nothing fancy here. This is ordinary regression that says how to work out the regression matrix if you use the basis that you can see on the screen.

You can see on the screen the penalties that are being applied to the coefficients associated with each basis function or hill. We achieve smoothness by penalising in rows and columns or along ages and years.

You can see on the screen what you get. This is a real attempt to get two-dimensional modelling going. The Lee-Carter model throws in the towel and says "I am not really interested in modelling. Do not even think of this as a model. This is a way forecasting and I think that it is a good way of forecasting."

The age period cohort model puts in quite a lot of structure and it may or may not hold. This is a real attempt to get a local model that will describe the mortality table, if you like, what is happening in that particular region of the table.

You can see that the functions that are fitted here are really quite markedly different at different ages.

You can see on the screen the forecasts that you get. All of this has been with the CMI assured lives data.

We are at the end of the talk, more or less, now. We have three methods of modelling and forecasting. The question is which is the best of the methods? Which one do I like? I cannot answer the question about which is the best method of forecasting but I can answer the question more or less about which is the best method to model. We have all five of our models arranged here. They are in the rank order. The 2-d method is the first. We then get the smooth age period cohort model, the smooth Lee-Carter model and then the original discrete models, age period cohort. Languishing in the last place, but do not forget, this is not a criticism, it is in last place as a model but it could easily be in first place as a forecasting method. If you read the Lee-Carter method original paper there is no suggestion in the paper that this should be used as a model. It has been taken up as a model. I do not think that Lee and Carter are very pleased about that. It is very much a forecasting tool. But you can evaluate it, nonetheless, as a model. This is all that you can do here since we do not have any future data, I am afraid.

The first column of numbers gives the deviance. This is simply the residual sum of squares for a generalised linear model. So it measures the goodness of fit, how close the model fits. That ranges from 6832 for the 2-d method. That is easily the best fitting model. It is not really a great surprise to me that that should be the case. It is the only model that makes a genuine attempt to respond locally to the two-dimensional behaviour in the table.

Following that, we have the age period cohort model as the next best fitting model. The discrete models are always going to fit better than their smooth counterpart. That will always be the case because the discrete model always can take local ups and downs to mimic how the data is doing. The smooth model cuts right through that.

So it never fits so well. That is the fitting - how well does it fit?

Then the next one we have is the trace or the degrees of freedom. That shows you how many parameters we are using to achieve this. The age period cohort model has a single parameter for every cohort. That is a huge number of parameters in its own right. It has getting on for 250 parameters here. The Lee-Carter model has two sets of individual age parameters and one set of individual year parameters giving it 196 parameters, more or less. The smooth models: we can calculate equivalent number of parameters, taking account of the smoothness. They use far, far fewer parameters. The Lee-Carter smooth in fact uses an equivalent number of parameters of about 30. The 2-d method is about 63.

So we then come to how do you rate these things? It is a bit difficult here. The discrete models are always going to score well because you have so many parameters. You have so much more room to do the fitting. We calculate a balancing function. This is called the Bayesian information criterion. It simply takes a view of how important the deviance in the trace is and combines these things to give an overall measure of fit that balances both goodness of fit, on the one hand, and the number of parameters you have used to achieve that fit on the other hand.

You can see on the slide where the rank order comes. The 2-d method is the best fitting -- the best model, in fact. Now we are talking about the best model as evaluated by this criterion. The worst fitting model is the Lee-Carter.

I emphasise again that we are not really here about modelling, we are here about forecasting.

Now I will show you a few slides that show you how these different methods work. We are now comparing the actual forecasts. You can see on the screen the slide for age 60. Here we have the Lee-Carter and the smooth Lee-Carter giving very, very similar forecasts. The 2-d method seems to think that mortality will not be nearly as low as that.

The slide you can now see is at age 70. The forecasts have come together quite closely here. This is because of the way the 2-d method responds locally to behaviour in the mortality table.

You can see on the screen the slide for age 80. The 2-d method is actually giving more dramatic mortality improvements than the Lee-Carter methods.

Unfortunately, I have not completed the work yet to continue the red lines. I hope to do that over the summer vacation.

You can see on the screen the slide showing all three ages shown on the same plot.

Now just a few concluding remarks. What I have tried to emphasise here is that this is about model choice. I have shown you three different kinds of model, and even within these three kinds of model I have shown you both discrete and smooth versions thereof. That is model choice. It is a big issue and it has a big impact on the forecast. We then have to choose the forecasting method. I have said just a little bit about that. Again, that is going to have a massive impact on the forecasting.

I have said nothing at all about parameter uncertainty. We can actually measure that in these models. But that is a different talk. The idea of parameter uncertainty is we have actually got the correct model but what are the parameters? We have to use the data to guess those and maybe you do not guess them very well.

Then we have stochastic uncertainty because individuals have a marked tendency to do what they will do.

You can see on the screen a few references if your work on pensions and annuities begins to pall. Perhaps you could look at some of these rather enticing academic papers. Thank you very much.

THE CHAIRMAN: Thank you, Iain. We will move straight onto Daniel Ryan, who is going to talk about another approach, Weibull. Daniel studied the medicine at the universities of Cambridge and London, before joining Watson Wyatt, in 1994. He was initially involved in the areas of critical illness and long-term care. During the last three and half years he has led a mortality and morbidity service for insurers and reinsurers in the UK. This service, which has eight current members, provides quarterly reports on research areas decided by the members.

He has recently been a co-author of two papers on the application of logistic and Weibull models to mortality experience for different countries from a human mortality database which were presented at the Society of Actuaries' Living to 100 and Beyond symposiums and the International Congress of Actuaries in Paris earlier this month.

Mortality Projection Methodologies:

Weibull

MR DANIEL RYAN: What I am talking about today is perhaps a slightly less sophisticated model than Professor Currie was talking about -- at least some of the models that he was talking about. What I am trying to put forward is a way of using a relatively less sophisticated model and perhaps adjusting the data as a method of projecting.

Effectively, I was going to put forward the idea that there are three different types of building a mortality model. The first one is possibly the one we are most familiar with, which is to use past trends, however you may define those past trends, in the aggregate mortality, and project them forward by some methodology without adjusting the data in any way.

A second methodology would be to adjust that past trend of information on aggregate mortality perhaps by changing in respect of some risk factor that you have more information on and you understand how mortality might vary in respect to that or by looking at alternative populations.

What I am talking about today is general population investigations, obviously moving to insured populations. That would be one such adjustment.

The third methodology is what I would say would be the more intuitive way of looking in terms of forecasting which is although we do not have any future information in terms of future mortality, we do have a lot more information on past mortality than perhaps we realise at first. We have the more detailed information relating to disease conditions underlying mortality, and while that data historically has been rather sparse, there are a number of different datasets which are now available, and increasingly available, which I know some of us are using to develop such models. I expect, as time goes on, and maybe we will hear more about this from some of the talks this afternoon, that that type of approach will be the one that is used.

But for this talk I am talking about the first approach with respect to a Weibull distribution. I will make a couple of comments how that might adjust in respect of the two methodologies at the end.

We heard from Professor Currie about the Gompertz model. It is a very simple model. Makeham adjusted that model to add in for the possibility of higher mortality not fitting in an exponential at younger ages. Kannisto then developed that. He was not so concerned about younger ages but was looking more at the older ages and the possibility for a plateau. You will see that if you choose the right figures you will actually get a plateau of one which is significantly higher than most commentators would suggest might be the plateau at old ages. Therefore there are other variants by Beard and Perks which enable you to have a plateau close to that which is being currently experienced.

In each case I would suggest that these models are fitting actual mortality experience in a specified calendar year primarily. You could fit them over the lifetime of an individual. You do have issues over the extinction of the cohort. Do you have the cohort all the way to death? Therefore are you fitting the entire function? Generally, these are primarily fitting in terms of a specified calendar year.

The Weibull distribution some of you may be familiar with; some of you will not. Effectively this was introduced by Waloddi -- I think that is a wonderful first name, he did have other first names such as Ernest. He was a Swedish engineer. Waloddi sounds much more impressive than the others that he had. That is probably why he used it for publishing.

The probability density function and a cumulative distribution function probably are not easily visible on the screen from the back of the hall. Let me point out the important features as I see them. You have within the probability density function a kappa and a lander. The kappa is the shape parameter which is effectively going to determine how the distribution lists to one side or the other. The scale parameter is going to determine the vertical axis.

You have a three parameter version as well which determines where the centre of mass of the distribution will be. But for this particular talk we have not had to use a three parameter version because a two parameter version fits adequately.

The two of these which are particularly important to us are the cumulative distribution function because that enables us to fit the data. But I will come to that in a second.

The reason that you may not have come across the Weibull distribution is primarily because it is not used for anything to do with human or other life. It is mainly used as a reliability distribution to model such things as strength of material, when you might get a failure when you are considering a mechanical or electronic device.

But more recently this Weibull distribution has been applied to consider different portions of the mortality experience. I say different portions. I mean at different age groups.

I will show a number of graphs showing Weibull. The very dark line, starting high on the left and curving down to the right is a Weibull with a shape parameter of less than one. The different graphs are using the same scale parameter but they are using different shape -- or sometimes it is referred to as slope.

So you can see by changing the shape parameter -- and I will go through which ones they are -- you can fundamentally change the way that the Weibull distribution looks. This is the probability density function that you are looking at.

Where it is less than one you have a very dark curve. That might be something that is appropriate for early failure in a component. But it also is not dissimilar to infant mortality -- perhaps less so now, but certainly in the past. As you increase the shape factor, then you see that the one which looks like a little hump and then fades away is 1.5 that the shape parameter is at. If you are at one that is a kind of random failure. The higher you go up to one to four, and, going back, the one you can now see is the one with the hump sliding away, is at 1.5. You will see the next one going across looks a little bit normal shaped. That is at 3.5.

So you are getting a pattern where failure is not expected at the beginning but is going to happen at some point in the lifetime and is going to be increasingly concentrated towards the end of the distribution. As you increase the shape, the skewness of the distribution increases, and so eventually here we have a B of eight and that relates to a fairly skewed distribution if you compare it to where you started, with B being less than one.

Why is this important? If you take the B up to these high levels, that is not that dissimilar to what mortality at older ages looks like. You start with a population of a certain age. You expect them to die off towards the end of the maximum of the individuals within that population. So this Weibull distribution has been increasingly used in a number of different papers. The first paper here was in respect of China by Vaupel. Vaupel will come to you later as being one of these proponents that you can keep on projecting forward past trends and not worrying too much because medical science, or various other factors, will enable you to achieve improvements because up until now we have not seen much convergence. What he did was he compared the Weibull to a number of other distributions, including Kannisto which we looked at earlier. He concluded that for the Chinese the Kannisto was actually a better fit but the Weibull was a better fit than certain other ones that he considered, such as Heligmann[?] and Pollard[?] or the quadratic.

Ozeki, in Japan, considered what he called a mixed Weibull distribution. In essence, you have already seen a mixed Weibull distribution because he had separate distributions for the younger ages where, effectively, the infant mortality he wanted to model differently, and he had a different distribution for the older ages. Therefore if you are considering only the older ages a single Weibull distribution might be something that you want to consider. This type of use of Weibull distribution he referred to as being used in some of the work on Japanese life tables. I have not personally read that in more detail.

Another paper by Weon looked at a Weibull distribution where effectively he had one distribution but the shape parameter changed with age, which is a similar concept to using a mixed Weibull distribution, although it might have slightly different shape in terms of how the shape parameter changes over time.

In essence, what he was suggesting was using the Weibull distribution, you might be able to model the mortality relatively well. I think that this is a fairly important point. What we are talking about is trying to model the mortality at this point and then later on we see whether this helps us in terms of forecasting mortality.

I have been involved in two papers as a co-author looking at the Weibull at a range of different countries taking data from primarily the human mortality database, which many of you will be familiar with. The first paper looked at the USA and Japan. The last paper looked at four European countries, although we cheated ever so slightly by separately considering West Germany and East Germany because of the way that the data had been presented. We have not a paper on England and Wales, so I thought that it would be appropriate to include a couple of statistics with respect to England and Wales. The principles that I am presenting after the slide you can now see on the screen are primarily based on information in those two papers.

The investigation period that we used was 1960 to 1999. That was entirely due to where there was commonality of information from the human mortality database when we did the work. There has been a more recent update of some of it, so you probably would be able to extend it to about 2002.

The countries that we considered were England and Wales, the USA, Japan, France and Italy. We also did some work on Spain and East and West Germany. The mortality measure that we are actually using to model is perhaps a slightly unusual one to you, but hopefully you will grasp it very quickly.

We are using probability of death from a starting age. Assuming you are alive, in this case, at age 50, then we are tracking exactly the probability in each future year beyond age 50 that you die. We are taking that from a vertical life table. You could alternatively do it tracking an individual through his lifetime, but then you would have a problem of cohorts perhaps not being extinct.

The way you could fit the Weibull are as follows. There are several ways that you could fit it. You could use a least squares approach in terms of graphically fitting it. The way that we chose to fit it was in using the method of percentiles. You may remember that, right back at the beginning, I talked about the cumulative distribution function being useful. You can use the actual mortality experience from that vertical life table to work out the probability of dying by a certain age as a cumulative function. If you choose 2%, I will say the 50th and the 95th -- we did -- you can then fit a Weibull curve that is consistent with those two points.

Then, by fitting Weibull curves in every single year that you considered -- and that is a fairly important point, we are fitting each curve to each calendar year separately and in different countries, obviously -- you can look at the parameters underlying the Weibull distribution and see whether there is any significant relationship in terms of time between the different parameter values. If there is a relationship, that may be a basis for projecting forward. That is not so much in terms of forecasting because I do not really believe it is forecasting, but in terms of projecting forward, and therefore being able to compare it to other methodologies of projecting forward mortality.

So we are going to start first of all with looking at the fit for a particular year. This slide shows England and Wales, and it is year 1999. The red line that you can see it is actual data. The probability of death is what we are measuring here, the probability of death from age 50. The blue line is expected. You may remember from the original Weibull graph that the Weibull looks like a normal but it then becomes skewed. So you expect the Weibull to come back down, and so the fact that the red line is starting to come down at the end is not really too much of concern to me because I know that the blue one is about to come down afterwards. So the fit between 50 and 80, demonstrated here, is pretty good.

The next graph that you can see on the screen is showing the fit for females. The fit may not be quite so wonderful here. But in essence we have a blue line which is tracking the red line relatively well.

I could spend the rest of today going through thousands of different slides showing the fits of every country in every single year. The slide you can see is a better way of demonstrating how the fit compares across years. This happens to be for USA males. The red squiggly line is the parameter value in respect of the shape parameter that is underlying each of the Weibull curves we fit to each of the individual calendar years from 1960 up to 1999, and the red line happens to be a linear line that goes through them which perhaps is better than certain other curves that we might fit to it. The blue line is in respect of the scale parameter.

You can see two things which are fairly important. First of all, the scale on the left-hand side is that relating to the shape parameter. You can see that we are varying from about 6.7 up to 8, which is kind of consistent with the type of values that we thought we would need to be able to model old age mortality. What is also interesting is the shape parameter is increasing relatively steadily over this period. There is some wiggling around it, but basically it is steadily increasing, which means that we are moving to a more skewed Weibull.

This is not really news to anyone, but it is showing that using the Weibull distribution, we are getting a relationship that we could project forward. How much credibility we can put to that is a separate point, but this is a methodology for projecting forward.

The slide now on the screen is trying to show you a situation where we only fit the data between 1960 and 1979 and then compare the fitted model which we project forward for the next 20 years with that which actually happened. In each case this is a life expectancy calculated as a vertical life table from age 50, and the scale on the left-hand side is the number of remaining years of life that you have. The blue line is England and Wales. The green line is the USA, and Japan is the red line. In each case the dotted line is the model continuing or projecting forward. The solid line is the actual experience, taking the whole population, that we are seeing.

There are a number of features coming up from this immediately. If your pattern of increasing life expectancy in the latter half was similar to the first half, then you would expect that the model would be quite good at projecting over that period. If, as was the case with the Japanese males, their rates of improvement fell off slightly in the most recent 20 years in comparison to the 20 years before that, then you see the model keeps projecting higher.

If you look at England and Wales, you can see where the improvements have been relatively better in the past 20 years. The model has not really caught up and is fundamentally lagging a little behind.

One feature that I would draw out very strongly, though, is if you look at the green and the blue line, at the beginning in 1960 the USA and England and Wales -- this is all general population, by the way -- were actually very close in terms of the life expectancy at age 50. By 1999 they were very close. The difference between the two seems to have been that the US had their improvements earlier. Obviously, there is a lot of discussion in terms of cohorts and how that fits in, but just viewed from that perspective we have a relatively similar improvement in the US which has faded off slightly in later years in comparison, as compared to England and Wales.

The slide now on the screen shows the equivalent comparison in respect of females. I would say that the Japanese one is one of those cases where you can quite clearly see that the pattern in terms of actual experience has been good before and after 1980 in terms of following a generally consistent pattern, and lo! and behold, the model does quite well at projecting the future. The model can only do as well at projecting for the future if they are using past trends as it is true to say that future mortality improvements are consistent with past trends.

In comparison, from the green line, which represents the US females, you can see that there has been quite a fall off in their improvements and therefore the model overshoots. The difference between the model and the actual in the females is much less than was the case for males for England and Wales.

We put all the different countries that we have modelled together, so we have England Wales, the USA, Japan in red, France in yellow and a rather fetching purple for the Italians. You effectively have here life expectancy measured from 2000. Once again there is a vertical life table from age 50. You can see that Japan is head and shoulders above the rest and there is a symmetry between the others in that because the way that the Weibull works, and because of the similarity of the mortality improvements in broad terms that are projected between the different countries because of the pattern being averaged over the 40 years, Japan remains above, but there is a generally steadily increasing rate of mortality improvement and therefore life expectancy. From memory, I think that we are talking about the order of about .1 to .15 years' increase in life expectancy per calendar year.

The females graph that you can see on the screen is on the same scale, so you can see women are doing much better. The Japanese once again are significantly above the others. There is some change between them, but fundamentally the others are showing similar patterns with perhaps France doing better than the rest.

Now, once again this is all just a projection. This is not forecasting. This is saying, "We have a model that we think fits quite well in the past. It seems to show a time-dependent pattern if you look at the Weibull distribution in different years. If you project it forward this is what will happen."

What does it actually mean in terms of the mortality improvements? You can see a graph on the screen which, on the left-hand side, is showing actual mortality improvements. I have smoothed them ever so slightly otherwise it would be almost impossible to read what they really were saying. On the right-hand side of the screen you can see the mortality improvements projected by the Weibull distribution being fitted in different years.

So you can see that there is a discontinuity. The reason that there is a discontinuity is because the Weibull has been fitted across all ages in the calendar year. It has not been fitted in terms of trying to pick up cohort patterns and projecting forward into the future.

The graph you can see is for England and Wales. You can see the blue segment represents the highest improvement in respect of males.

The slide you can now see is the equivalent graph for England and Wales females. You can see a slightly lower level of improvement. It is still a diagonal. The rate of improvements under the Weibull distribution is lower than that which we saw under the males.

The slide is now showing the US males. The Weibull improvements are actually quite similar between the US and England and Wales. In part, that is because the mortality change over the 40 years we have considered and the Weibull fitted over 40 years is quite similar between England and Wales and the USA.

So, in assuming a different rate of mortality improvements between England, Wales and the US you are fundamentally saying that there is something that is going to happen that will be different between the two countries. I will come back to that point. That is an assumption. That it has been the case in the past does not necessarily mean it will be so in the future.

Two features which are still rather important to bring out from the Weibull comparison with some of the mortality improvements that you have seen elsewhere are as follows. Once you get past the cohort period in the medium cohort, or whichever cohort you are looking at, or if you just look at CMI 17, you have a rate of mortality improvements that reduces as you increase with calendar year and as you increase with attained age. That is less of the case with the P-splines or Lee-Carter where it is flatter -- flatter in the case of Lee-Carter and with the P-splines you get a continuation. But here you are seeing a pattern whereby basically the highest improvements in the Weibull section are at the bottom right and the lowest are at the top left. As calendar year increases, your attained age has higher mortality improvements. As the Weibull shifts more to the right, more work has to be done at the older ages to sustain the level of improvements in terms of life expectancy which the Weibull is projecting and therefore the mortality improvements are increasing.

This slide shows the USA female population mortality improvements. Once again, it is useful to look at the USA in terms of females and males. This has been remarked on. You do not see diagonal patterns; you see vertical patterns in that the improvements seem to happen in a block rather than being spread out or associated with a particular cohort.

If we are going to project forward we do, at some point, have to make a subjective decision. The choice for all of us is where that subjective decision happens. The subjective decision could be at the level at which you think the past trends cannot be projected further forward into the future. As I said, there is this view by the people like Vaupel and Oeppen, and you have probably seen this in other presentations, which is showing a comparison of life expectancy in various industrialised countries, and how it has changed over this period of 160 years. The one with the highest life expectancy has continued to break through these barriers that various governments and the UN thought there were and therefore, Vaupel and Oeppen would suggest, will continue to happen. Others, such as Olshansky, have pointed out that that is wonderful but you have to justify that in some way. You have come up with some level of improvement which is consistent with a particular change in a risk factor, a particular introduction of a treatment or some other effect on the morbidity underlying the mortality.

The way that you might adjust what is effectively a very simple Weibull distribution to be more appropriate for projection and therefore ultimately for forecasting is first of all, most obviously, you may apply this distribution directly to insured population experience. That is either in terms of the CMI's own experience or you may, if you have sufficient information, be able to do it with your own offices where you are part of a life office.

Secondly, you may want to adjust the data to the extent that there are particular factors, perhaps smoking, where you have a view that in the past trends so much of the mortality was associated with changes in the prevalence of smoking and the level of dosage of smoking that was inhaled, and how that might be different in the future and therefore adjust the mortality improvements that you have seen the past to be more appropriate to the future.

Finally, and this takes us back to our original point, there is a spectrum where you continually adjust the past trends to take into consideration information that you have available to you. That may be just looking at particular risk factors that you have a decent model for or that may be in terms of adjusting between the differences between the population that you are interested in and the population where you have data. You may be able to do that through cohort studies; you may be able to do that through introductions of new treatments. At some point you are shifting the subjectivity from at what point do you assume the past trends run out to at what point do you bring in subjectivity from other professions? I understand, quite reasonably, that actuaries are concerned about involving subjectivity outside of their profession. But it is an area which, I think increasingly, we will move towards, and it is probably better if we embrace it ourselves rather than other professions coming up with models which we have to accept.

THE CHAIRMAN: I should like to introduce Phillip Olivier and Andrew Smith. Phillip is the senior manager in the life actuarial practice at Deloitte, which he joined in 2004, since when he has been involved in projects relating to realistic balance sheets, ICAs, shareholder value, bulk purchase annuities and product pricing.

He has assisted clients with the implementation of stochastic totality models, mainly focusing on insurance risk and annuity portfolios.

Mr Andrew Smith, possibly needs no introduction, but I have been told to introduce him as an actuarial student who works for Phillip!

Mortality Projection Methodologies:

Olivier-Smith Method

MR PHILLIP OLIVIER: Basically, for those of you who were wondering, Andrew does not really work for me, but what we are going to do here today is a bit of role-play. I am going to be the clever actuary working for a life company and Andrew is going to be one of my actuarial students.

Andrew, I have an interesting actuarial problem that I should like you to assist me with, hoping that this will give you some exercise and maybe argue with those exams. What I should like you to do for me, please, is price a stop loss reinsurance contract where the direct insurance company basically sells 100 annuity contracts to a cohort of lives aged 65. My calculations, using my own assumptions, show me that the total annuity cash flows would be £2.12 million. The stop loss reinsurance contract will start paying as soon as the total annuity payments exceed £ 2½ m. Can you please go away and calculate these the stop loss premium for me?

MR ANDREW SMITH: I was a bit worried about that task, Phillip, because I know that that £2.12 million has come from our own mortality assumptions. I dug around where those have come from and it seems that some of it is from UK experience and the some of it is based on some international experience that we thought was relevant. We have adjusted it a bit to allow for our particular sector of the market that we have. We have also adjusted it a bit by talking to some gerontologists with some views on what might happen in the future. I thought, that is great, we have a really good mortality table. We have asked all the leading people in the industry how to produce that.

I went to the CMI website and I tried to download their stochastic projection models. I ran into a problem. The problem was that they are models they project starting from their own data set. If I want to get a projection where the average corresponds to where we think our mortality table is, it seems we have to reverse-engineer an invented data set such that when you put it through the CMI sausage machine, the projections that come out on average look like where we started from. So I found that pretty exasperating that I was stuck with their base tables as well as whenever I wanted to use their projections. So I wonder whether you might be able to give me some guidance how I can get around that.

MR OLIVIER: Andrew, I always say if you want something done properly, you should do it yourself, so I have gone away and I have basically devised a method that enables me to use my initial estimate of future mortality as an input into my stochastic mortality model. That initial assumption that you used can include expected future trends and the impact of cohorts, as you see fit. I have basically used this model. I have produced some results.

We can see some results on the slide on the screen. What is shown is the tail of the distribution, the total annuity outgo and the different assumed levels of volatility, using this model.

MR SMITH: Phillip, I am a bit worried about that. What you have projected here is the total number of deaths. We might know that if we wait a very long time until the last person has died. But when we do our pricing of these products, we try to project financial statements, tax, and things like that. So what I need to do there is to project the future mortality tables. I have had a few ideas of how we might project the mortality tables.

The first idea that I had is start with our best estimate and then use the stochastic projection of number of lives and assume that whenever we have to draw a future balance sheet we use today's mortality basis.

That was not quite right because if mortality improves and some statins start working, then we will revise our tables. So then I thought again. I thought maybe in each simulation we could look at the pattern of future deaths and use that to reverse-engineer a mortality table for a projection on that simulation.

I realised that that was not quite right, either, because that would assume that in one year's time we then from then on have perfect knowledge of how many people will die. So I am at a bit of a loss how to project my mortality table that I am going to use in one year's time to value those liabilities.

One of the reasons that I am worried is because I remember when we had the AIDS scare we actually strengthened our mortality bases quite heavily. We took a hit to income and then we ended up releasing most of it again because AIDS was nothing like as bad as we feared. So I know that, to some extent, tables can have a life of their own, even if it is not in the experienced number of lives.

Can you give me any guidance as to how we might go about that?

MR OLIVIER: What you could do is construct a stochastic mortality model which enables you to model both expected future mortality experience as well as valuation tables. What you can do is start with an initial mortality curve, including your best estimate of future improvements.

As time passes, you can revise each survival probability according to a formula where new survival probability is your previous survival probability to the power of a certain deterioration factor. Your new survival probability is the one calculated at the current valuation date. The old one is the one at the previous valuation date. The deterioration factor is basically a stochastic quantity. It follows a gamma distribution with a mean of 1 and a certain variance. The reason why you need a mean of 1 is because you started off with a mortality table which already contained future improvements. So what you want this model to do is to pick up any deviations from your initial expectation.

What you can then do with this model is apply it inductively for multiple time horizons.

MR SMITH: Phillip, I have tried to use this model, actually. Again, I got stuck. I will tell you the problem that I had.

Using this model, I projected all these simulations of future forces of mortality and I found that when I averaged the force of mortality across the simulations I got back to our initial table that our analysts had come up with. I thought that was pretty good.

Then I projected the number of survivors. I cannot quite figure this out. I found that more people were surviving on average from these simulations than I started off with for my mortality table.

I do not understand that because my forces of mortality, on average, were bang on, yet the number of survivors from my simulations always seemed to be a bit too high. I tried with the different random number and was always still too high. I never got it coming out too low.

So I do not think that it is sampling error. But I am a bit baffled. How can it be that if I calibrate my force of mortality to be right on average, the number of survivors is overstated compared to what our expert is telling me?

MR OLIVIER: That is a very interesting point, Andrew. Basically, you have to choose whether you want to get your force of mortality correct or whether you want to get your projected number of lives, and therefore the cash flows produced by this model, accurate.

The reason why this happens is if you get the force of mortality correct, in low mortality scenarios that low force of mortality will apply to more lives. Because when you perform lots of stochastic simulations you are effectively calculating the weighted average, your mean from your stochastic simulations is going to exceed the projected number of lives produced by deterministic projection.

So in order to fix that, one needs to introduce a bias correction factor into this stochastic mortality model to ensure that your cash flows are not biased. Clearly the work that we as actuaries do is more concerned with the cash flows than with the mean force mortality.

So, let us introduce some notation to explain how this will work. We started with homogenous cohorts of lives. Let L(t) denote the number of survivors from that cohort at time t. Let L(s, t) denote the expectation under s of the number of survivors at time t. That is conditional expectation basically that takes into account all information available at time s.

The slide you can see on the screen is a fairly ugly slide. I am going to have a go at trying to explain this as quickly as possible. I might not get it 100%. The top line says the actual number of survivors at time s is distributed binomially where the sample size is the actual number of survivors at the previous valuation date, so it is time s minus 1, and where the probability of surviving from time s minus 1 to s minus 1 is calculated as the conditional expectation at the previous valuation date, so time s minus 1. The starting point there is the binomial model.

The next line down says that the probability of surviving from time t minus 1 to time t, as measured at valuation date, s, is the same probability measured at the previous valuation date to the power of this gamma innovation which I have talked about before. That is then multiplied by a bias correction factor. This is the bias correction factor which you need to make sure that your cash flows are unbiased.

The bias correction factor is what you need in order to ensure that this life function of yours is a martingale in s, and you get a fairly ugly formula. But if you use that formula then the result is that your cash flows are unbiased.

MR SMITH: I am worried, Phillip. Have you not been a bit simplistic there? Your formula is there. I think that is all for one cohort. So that assumes that I have got lots of people all the same age, all the same smoker, non-smoker status, similar health conditions and similar levels of annuity, that they are being paid. So I am not really picking up the diversity of the population, am I? When I have different mortality tables for different sectors of the population, I am a bit stuck as to how to use this model. Can you give me some ideas?

MR OLIVIER: Andrew, fortunately, the solution to that is quite simple because the model uses as an initial input your mortality tables, including expected future improvements. All you need to do, basically, is to use different tables for your different cohorts of lives. As long as you use different mortality tables by age, gender and smoking status, etc, you can project separately for each of these cohorts.

However, the way that the model is constructed at the moment it will use the same stochastic deterioration factor to all of those cohorts. What that means is mortality will improve by more than expected originally or deteriorate by more expected originally for all these cohorts together. So for everyone at once. It is probably worth mentioning that one could extend the model using principal components analysis to do something more sophisticated. That is the way that it is constructed at the moment.

MR SMITH: Thanks for that. I think I can project the mortality. There is one last thing that is worrying me. We are not only exposed to mortality risk, and in particular we have put on this interest rate hedging program, and the interest rate hedging program has assumed that we know what the cash flows are going to be, so we have our mortality table, we have projected our cash flows. I am concerned that if mortality assumptions are wrong then our interest rate hedge will not work anymore. So if mortality does not improve as fast as we expect, that will shorten the duration of the liabilities and then we will be potentially exposed to a rise in interest rates. How can we consider those together?

MR OLIVIER: Another good question, Andrew. The techniques that enable you to identify those combinations of stresses will have the most severe impact on your balance sheet. Those would be the ones that you would want to focus on if you devise such hedging strategies.

What you can do is you can start with your base case. That is the blue line in the middle of the diagram. You can perform a few one directional stresses, which we call naive stresses, so you can stress mortality in one direction and then the other, and you can do the same for interest rates.

The technique enables you to determine a collection of stresses which all then have the same probability of occurring. That is the purple ellipse that you can see on the slide. We call that the likely locus. The technique then also enables you to determine which stress on that ellipse will hurt your balance sheet most. We call that the least likely solvent event. That is the event on that ellipse which is going to hurt you most on your balance sheet.

That would be the stress that you would want to focus on, Andrew, when you put a hedge in place.

It is also worth mentioning that it also enables you to determine the event that is most likely to make you go bust. So if you look at the red line on the slide, all of those are combinations that will make your company go bust. The most likely ruin event, the red dot, is the one that is most likely to happen.

MR SMITH: I think that I get it now. The reason that that red line slopes is that is saying that if mortality improves and I go down this vertical axis, I have lower mortality, so therefore I have longer liabilities, and that is why I am exposed to a fall in interest rates in that scenario.

MR OLIVIER: That is correct.

MR SMITH: So would it make sense if I lengthen my assets? If I held assets that were more like the liabilities in the stressed mortality scenario? Then I might flatten out that red line, might I not? Might that mean I need less capital?

MR OLIVIER: Yes, correct.

MR SMITH: That is really interesting, thanks.

MR OLIVIER: Briefly, to summarise, we have looked at a few potential issues with using stochastic mortality models. One might be just from a logistical perspective. You might not be able to use your own mean assumptions. Using the Olivier-Smith model you can use your initial curves as an input. You need to project balance sheets if you do solvency work. This model allows you to model experience cash flows as well as mortality tables. You might have a bias on your cash flows. A biased correction factor can be introduced to assist with that.

You might need to model multiple cohorts. This technique enables you to use different assumptions for your different cohorts.

Finally, you need to look at all your risks together. Techniques exist to let you perform your stress tests a bit smarter and determine which of those you need to focus on.

It is probably worth mentioning there is no need to look at mortality and interest rates, you can look at any combination of any number of risks at the same time.

Questions

THE CHAIRMAN: Thank you to Philip and Andrew -- your junior actuarial student -- too!

The plan when we put the programme together was that we now had a meaningful amount of time for discussion and debate. There has been a wonderful content of material to try to take in. We are now opening up the seminar to the floor for questions, discussion points, issues arising from any of Iain, Daniel or Phillip’s presentations. May we take the first question, please?

MR MIKE HARRISON (Mercer): People do not generally get killed by B-splines and gamma deterioration. They get killed by things like cancer and heart attacks. But there is only passing reference in Daniel's slides to the fact that we need to take into account things that we know, like medical improvements, and what mortality is doing in other countries. So I am a little bit concerned that we are missing a trick as a profession and that perhaps we have missed a trick in the past. I wondered if you could comment on that?

THE CHAIRMAN: There is an awful lot of people up here who could comment. I am hoping that Richard Willets will bring in that sort of reality check after lunch.

I think with the models that I have used and I have experimented with, the key thing is yes you do have to have a view. There are various ways to fit the models and therefore you do need a view of where you are trying to head. The subjectivity does come in in the model fitting.

PROF CURRIE: This is the whole business of cause specific death and how that can be incorporated into the forecasting. My own view, which I tried to lay out at the start, was that all of these different effects are rolled up into one general effect which then acts on the whole mortality table. It is this which is producing the improvements that we see over time. It is what we expect these as yet undiscovered medical advances will allow the mortality table to continue to improve in the same fashion as it has in the past.

I am always slightly uneasy about cause specific when I identify a particular disease. It is very easy to imagine that you could miss something else. It is certainly the case that if you have mortality on the data that is divided into separate categories, and they each have their own mortality table, then you should certainly model separately like that. I do not think that you are really talking about that. You are imagining a supermodel that incorporates into the models that I was presenting something about particular causes. That is a bit more problematic...

[TAPE CHANGED AT THIS POINT]

...THE CHAIRMAN:... and we tend to think of morbidity as something that has to be dealt with in critical illness and mortality as something that has to be dealt with in annuitant mortality or term assurance separately. I think a slightly more joined up way of looking at it is not possible through the way that we collect insurance data but is possible through the way that various population data are collected. Maybe that is what Richard will talk about this afternoon. By using that type of cohort information, where we are tracking individuals with particular diseases, and seeing not only when they die from that particular disease how the presence or absence of second disease affects that point, we can get a better understanding of how, for example, the reduction of the particular disease may really impact when you allow for the interaction between diseases.

It is only through paying attention to those kinds of population cohort studies that I think that we will be able to answer, or start to answer, those type of questions.

MR SMITH: I think there is a danger of doing more complicated models than are needed. If you are paying out on life assurance, you pay out if somebody dies. You do not generally pay out a different amount for cause of death. Similarly, for annuities.

The way that I look at it is you can build a model of general mortality patterns and then if you have ended up by saying, "My capital is being determined by reference to this extreme mortality event" that is the time to do your reality check, and say "Can I paint a picture of how that mortality improvement could happen for the population?" It does not necessarily mean that you have to build up your forecast from the bottom up.

I give you an example. I can look at how long it takes me to get to work and the delay to my train. I could say "I cannot possibly analyse railway delays unless I analyse them by course." I might consider trackside fires, staff shortages, points failures, snowfalls and strikes, or whatever. But do I actually get a better forecast of my delay tomorrow by doing that breakdown of cause of delay? It is by no means clear that I do.

So to my mind there is certainly an awful lot that can be learnt simply by drilling down more deeply into the data that we have on numbers of deaths.

MR [?] (Hannover Life Re): I have a couple of questions. The first one I should like to ask of Iain is this. The penalised spline method, the actual B-splines themselves, you said that it is invariant to the number of splines used, provided that you have a redundant number of splines. Is there anything magic about the formula for the actual spline itself? Could you use any old function? Could you use a wider spline that covers a wider age range? Do those things make sense to look at? Or do they not really add value to the whole footing process?

PROF CURRIE: I showed two examples, you remember, one where I had the knots, the distance, which determined the width of the spline. I had knots where they were spaced at two years and nine months, and I also had knots where they were spaced at 11 years. In your way of thinking of it, they would be big, fat splines. If they were two years and nine months they would be skinny splines. It is depending how local the functions are. They are very local and very responsive.

MR [?]: I had not appreciated that the splines changed shape, depending on the number of splines that you used.

PROF CURRIE: They expand and contract.

MR [?]: One other comment was on Daniel's presentation about using the Weibull survival function to project mortality and possibly model existing mortality rates.

There is a famous Russian couple that a person has to mention at everyone of these conferences, called Gavrilov and Gavrilover[?]. They have come up with a very nice mathematical theory about the factors that drive mortality. What you said about the Weibull model underpins that, really. The Weibull model is a method for modelling failure rates of highly engineered combinations of systems.

The logic here is that Gompertz is better at modelling failure rates of systems with lots of redundancy. So you have a system that can take a knock, that can have a heart attack, and there will still be surviving heart tissue, and so on and so forth. They have come up with a very simple, very elegant, formula that combines the Weibull survival function with the Gompertz survival function.

They call it the binomial law of mortality. It is just a hazard rate model with three parameters. I have messed around putting it to various mortality databases and it works beautifully. You can describe tonnes and tonnes of things with it extremely flexibly. It captures late life plateaux, if you believe in that sort of thing, and all that sort of nonsense, so that is quite useful.

I have a question for Andrew: I did not quite follow the logic behind the bias. How does that work? You say you have low mortality rates. Is the reasoning that when mortality rates are unusually low those are the situations when you have more survivors than you expect and it contributes disproportionately to the weight of average?

MR SMITH: Yes, that is exactly right. You are looking at a simulation environment where you are starting from a table today but you do not know what will happen in the future. Imagine if you have two possible scenarios. One of them where the force of mortality was always 10%; the other of which the force of mortality was always 20%.

If you take the average of those two scenarios, the combined force of mortality actually tends the 10%. That is because the 10% of mortality cohort lives longer and ultimately makes up virtually all the population. So that is a problem for calibrating simulations because you would like to calibrate so that your initial annuity prices are, on average, what your annuity prices are.

MR [?]: It is a novel insight, really, because it is not the sort of thing that I would automatically think about when I am doing a stochastic projection.

MR SMITH: You are right. It is not a problem that I have seen described elsewhere in the literature. Much of the literature is to do with getting forecast models. The commercial issues that you encounter when you try to use them for commercial problems we are just starting to see emerge.

MR [?]: Then one final comment on trying to predict the essentially unpredictable and looking at cause of death and trying to tie all these things into predicting future mortality improvements. It is very complicated and I think that the title of what we are talking about is right, it is the highly uncertain.

A strong argument that has emerged recently, based on some fairly solid research by reputable researchers, is that one of the big drivers of extremely late life mortality, the point where you have entered the stage of life where you have very little redundancy left and any particular problem that comes up is likely to kill you, is they found a big driver of that is just early life conditions.

In the first couple of years of your life, depending on how many infectious diseases you are exposed to, depending on all the problems you had, nutrition, and so on, and so forth, early life conditions are a large driver, they have concluded, to those late life mortality levels.

If you extend that idea and say "Look at how early life conditions have improved" well, they stopped improving when? In the sixties? I am just being blasé when I say that. I am not sure what the number is. But some people who have looked at this would then look at the future mortality improvement rates for the extremely old and say, “A big component of those improvement rates is set to continue until 2010, but early life conditions have not really improved that much, so that component of future improvement we will still witness for the next 10 or 20 years, but it is going to fall away and then we will rely on the future unseen things."

There are so many things in there that I think an aggregate approach is probably the only way to do this.

MR MATTHEW EDWARDS: I think it is a question mainly for Iain Currie, although there may be ramifications for the other speakers about the choice of the smoothing parameter lander which seems to be particularly a feature of yours, although presumably there are equivalents in all the other methods how you end up smoothing it. It seems to me that, for reasons of time, you skipped over that. That seemed to be a fundamental part. I just wondered to what extent there is subjectivity, as opposed to objectivity, in choosing that.

I was wondering in that context, also, if I am right in thinking that you can get part of the way of setting it scientifically by splitting your data, so you use part of the data for setting the other parameters and then perhaps choose a smoothing parameter on the other part of the data and then look at the fit on that, or something along those lines.

I was wondering if you could you speak a bit about how the smoothing parameter would, in reality, be chosen.

THE CHAIRMAN: Is this in terms of the ranking of the different models? Basically that involved, presumably, some balance between smoothness and goodness of fit.

MR EDWARDS: With particular reference to the penalty likelihood. We have an interesting formula with a sort of half lander. I suppose another question could be why you have a factor of a half in there.

PROFESSOR CURRIE: The half is there so when you differentiate a quadratic you get a two, and the two will cancel the half. If you want to take the half out, the lander would become twice what it was.

Regarding your question, yes, this is a huge issue, of course. People have thought a great deal about how you balance goodness of fit, on the one hand, and smoothness on the other. There are various "scientific" approaches to this. If you adopt a Bayesian approach, you can show that the minimum errors in the future forecasts lead to a particular formula for the value of lander. That is the one that I have used. But you do not have to agree with that particular approach. There are other approaches.

There is something called the [?] Information Criterion, which tries to measure information content as a compromise between the two things. That will give you another value of lander. It will give you a much rougher function.

You can also use what you were referring to, which is fitting on some of the data and checking to see how it does at predicting the other part of the data. That is called cross-validation. You can certainly choose lander using cross-validation techniques. That is another approach.

Which of these is the best way of doing it? Who knows? Subjectivity comes in there. I chose the Bayesian information criterion because it does tend to give rather heavy smoothing and my attitude was that if you were trying to forecast then perhaps heavy smoothing was not such a bad idea. You were more likely to get a stable forecast with heavier smoothing than with lighter smoothing.

It is something that I did not have in my talk, but I was talking to people before we started and somebody was saying that they have been using the 2-d method on a number of countries and were commenting that they were a bit disappointed on how it worked in certain countries.

I should just like to say a little bit about that because it is probably something that a few of you have come across.

The important thing about the 2-d method is that it will give you a good forecast only if there is a strong signal in the time direction. If the signal in the time direction is not strong, then the method will not work well as a forecast. There are some countries, for reasons I do not understand and I have not investigated, where the signal in the age direction much dominates the signal in the time direction. The forecasts used for these countries are not at all satisfactory, the reason being if the signal in the age direction is very strong then you essentially just get a linear forecast. The thing linearises across age, and you really do not want that.

So you do have to look at the model, and say "Can I sensibly use this model for forecasting?"

MR SMITH: Ian has talked about the Bayesian information criterion as being one way of essentially looking at data to see how much wiggliness there is and therefore see how much you need to smooth it.

Iain's observation was that that would tend to do quite heavy smoothing. From some of the data I have seen I suspect that actually you still need to smooth it even some more. The reason for that is that quite a lot of the jumps in the data can be due to funny things that would not recur in the future.

I will give you an example. When you are looking at population data, there is often debate about how you collect census information because not everybody responds to census forms and the people who respond are not evenly distributed among social groups, ethnic groups, and so on.

So what you find is that the central government departments will decide "This year we are going to gross up the census data in this way". Then, three years later, they will change their minds.

That flows through into an apparent change in rates of mortality. But all that is actually going on is the ad hoc adjustments to the exposed to risk have been changed at the whim of somebody in the statistical office.

If you are fitting splines, you will end up coming up with an idea of having to have fairly narrow splines in order to capture all those changes. But if you are actually using that population data with some adjustment to allow for assured lives, you have a much better idea of your exposure to risk and you are not going to start grossing up census data to work out your exposed to risk, and therefore the hiccup changes you get in the past data really are not relevant to projecting at all. So I would argue that actually many of the mechanical methods for finding the amount of smoothing tend to smooth not enough compared to what can probably be justified.

THE CHAIRMAN: I am going to draw matters to a close now, slightly later than we intended. May I ask you to thank the speakers? [Applause]

Mortality Projection Methodologies:

Recent Developments in Mortality

MR RICHARD WILLETS:...

[FOLLOWING LONG BLANK PERIOD ON TAPE]:

... research, and talking to medical experts in this area, is that it is entirely feasible that we are going to see, and continue to see, very large reductions in heart disease mortality in the near future, and the debate really is how we can best achieve those reductions as quickly as possible.

Moving on to lung cancer, what the graph you can see on the screen actually shows is over rolling periods of ten years the proportion of the total improvements in mortality for men in four different age groups that are due to lung cancer as a proportion of the aggregate. What you can essentially pick up from this graph is that the proportion due to lung cancer is actually reduced from its peak from the ten years to 1994, so now it is ten years to 2004. Lung cancer is a less significant cause of death as a proportion of the total improvements.

Lung cancer, in some ways, is a very interesting cause of death from a mortality point view, because it is very strongly linked to lifetime cigarette consumption, and you can do some very powerful modelling of lung cancer mortality by looking at individuals' average consumption of cigarettes over their lifetime. This is shown on the two graphs that you can see on the slide. The one on the left shows female cumulative consumption of cigarettes over their lifetime. Each line represents a different age group. On the bottom axis you have year of birth. It shows that lifetime cumulative consumption actually peaked for females born about 1925.

Then on the right-hand side of the slide you have the rates of lung cancer mortality. Again, each line represents a different age group. You have year of birth along the bottom of the graph. You can see very clearly that lung cancer mortality, also peaks for each age group for women born about 1925.

In fact, the ONS, in a publication about ten years ago, "Health of Adult Britain" described female lung cancer mortality as being "an almost perfect example of a cohort effect". So year of birth is really the prime determinant over time of lung cancer rates.

Again, I fitted my age period cohort model to lung cancer trends and that did indeed suggest that by far the most dominating factor was year of birth. It also suggested that calendar year was a significant factor as well, not as significant, and that the pace of improvement by calendar year actually peaked in 1990 to 1995, so the early nineties.

The reason why it is interesting, and it is interesting in itself because lung cancer is a significant cause of death, is trends in lung cancer mortality are sometimes used as an indirect indicator for decomposing aggregate mortality rates into parts that are attributable to smoking and parts that are attributable to other factors.

This would suggest, along with the model of heart disease mortality, as illustrated earlier on, that the impact of smoking was greatest on mortality improvements in the early nineties and that since then the impact of smoking has somewhat worn off.

Many people have been talking quite recently about the potential impact of smoking bans, bans on smoking in public places in various countries. In trying to gauge the impact of that, it is certainly worth bearing in mind if you are considering the impact that that might have on the mortality of annuitants or pensioners who are receiving their pensions now.

Cigarette smoking prevalence is actually very low at advanced ages. On average only about 8% of people in England over the age of 75 are reported to smoke cigarettes. As well, epidemiological studies show the benefit of giving up smoking reduces significantly with increasing age.

A paper a couple of years ago suggested that a 30 year old smoker, if they were to give up cigarettes, could save approximately ten years of life, but a smoker aged 60 would gain only three years. That obviously reduces quite drastically with increasing age.

I would imagine that the potential impact of the introduction of smoking bans, no matter what the impact on smoking prevalence will be, at older ages is going to be relatively minor -- probably not a material effect. Where it may have more impact is on younger generations, people now in their thirties, forties and fifties.

Considering other cancers, you can see on the screen the same graph as I showed for lung cancer. It is the proportion of the total change that is due to changes in the mortality from other cancers for four different age groups.

You can see the reason why you are getting negative numbers on the left-hand side of the graph is that cancer mortality was actually worsening during the 1980s for people in these age groups. But steadily over the last 10, 15 or 20 years, we have seen a reversal of that trend. So now changes in cancer mortality are accounting for perhaps 10% of the total mortality change.

Why has this happened? Here is one instance in which improvements have been mainly driven by medical advances. In some cases the patterns are very different for different causes of cancer, but in some cases recorded incidence rates have increased whereas mortality rates have reduced. There can obviously be lots of confounding effects which influence incidence rates, such as the introduction of screening programmes, and so on; early detection of cancers. Breast cancer is one example in which that has happened. But it also shows very clearly how you have a very divergent trend between incidents and mortality. This is the change in female breast cancer mortality and incidents over an 11 year period for various age groups.

Generally speaking, for females in their fifties and sixties you have had an increase in incidents of perhaps 10% or 20% over that 11 year period. But the change in mortality has been a reduction of about 30% for most age groups.

This is also consistent with changes in five-year survival probabilities for female breast cancer diagnosed in women in different periods of time. You can see from the slide on the screen women diagnosed early nineties had a survival rate of about two-thirds, on average, whereas for women diagnosed in 2001-03, a five-year survival is projected to be about 80% for all ages.

So you can see the reason for this improvement in mortality in breast cancer is improvements in the screening and treatment of women with breast cancer.

We are just about seeing on the radar screen mortality improvements for cancers beginning to influence the overall mortality trends. If you gauge people's views about what might happen in the future some people suggest, looking 50 years out, the way we look at cancer may be completely different to the way that we look at it now. People suggest that it might be as controllable as diabetes, or a large percentage of cancers will be controllable at this point in time.

If that were to be a scenario which happens, then the improvement rates that we are seeing that the moment from cancers are going to have to accelerate very rapidly in order for this to become a reality.

In summary, what I draw from exploring the trends in these three causes of death is a few points. First of all, the relative importance of cigarette smoking as a driving force of mortality improvement has diminished since the early 1990s. But this has not lead to the predicted reduction in the aggregate pace of improvement. You might say "Who predicted that mortality reductions would diminish because of trends in smoking?" I predicted that, for one. In the nineties I would have said that that was the most likely to happen, the improvements due to smoking would run out of steam and therefore aggregate mortality rates would reduce. I think that I was right in saying that the mortality improvements due to smoking have reduced. The only thing is that those improvements have been replaced by improvements due to another cause, which I think primarily is medical advances, both in relation to circulatory diseases and in relation to cancer. So medical advances, I think, are playing an increasingly important role in driving mortality change.

I suppose this is almost an aside, or a further point. It is quite interesting, if you go back to my age period cohort model again but this time consider the mortality trend for all causes of death but excluding heart disease. Essentially, what you have on the slide on the screen is this period component shows the rate of improvement stripping out the cohort effect, stripping out the really rapid improvements that we have seen in heart disease mortality.

What have you got left? For males, and you see the same sort of trend for females as well, I think what this indicates is more or less over the past 35 years or so, we have seen a fairly even rate of improvement once those factors have been taken out. There are slight peaks and troughs but on average that rate of improvement, that period component, has averaged 1.7%.

I said that I was not going to talk about projection methodologies, but I cannot really help myself. The slide that you can now see shows some projected improvements implied by the medium cohort basis, which I think is still quite widely used in practice. These are annual rates of improvement for somebody aged 75 over time projected by that basis. You can see that they are quite high at the moment. In about five years' time they are projected to have reduced very sharply. A characteristic of this projection basis is that over time the rate of improvement falls down to a much lower level than we have seen, I would say, over the past 30 or 40 years, down to about ½% per annum once you get out about 50 years or so.

My own personal view on that particular matter: the issue of using projection bases which have seen very low rates of improvement in the long run. I stress that this is my own personal view, not the view of the CMI or anybody else that I have ever worked for, met or known.

In most instances, I would say mortality projections which generate future improvements far below the rates we have seen over the long-term, such as the medium cohort basis, I do not think are suitable for generating best estimate assumptions for reporting embedded value or pension fund reporting, primarily because (a) they run counter to the trend that we have seen over a long period of time; and (b) they also seem to run counter to a forward looking approach of mortality change, looking at the potential drivers of mortality change, talking to medical experts and considering what might happen to different causes of death.

Questions

THE CHAIRMAN: Thank you very much, Richard. A fascinating presentation, as we were expecting. Questions, please?

[?]: To what extent is it the case that when you get rid of one cause of death, another one comes along to replace it? For example, because it has been in talked about, but I would not want you just to focus on it, there is the issue of obesity.

MR WILLETS: The issue of the interdependence of cause of death is very significant; the older you get, the more difficult that issue actually becomes. One thing I would say is you can actually gauge quite a lot from looking at the recent trends because we have seen very big reductions in certain causes of death. You can see what happens when that happens. You can see what happens when we have massive reductions in heart disease mortality. Does that lead to an increase in stroke mortality, because the two diseases have quite similar drivers? Has that happened? The answer is "no". You see reductions in stroke mortality as well. Have you seen increases in cancer mortality as well? Probably not. They have reduced, too.

It is a very difficult problem to model correctly. In fact, some people would argue that it is impossible to analyse all the interdependences in the different causes of death.

I suppose, turning to obesity, that has created a lot of new stories in the last couple of years. I think potentially it could have quite a significant effect on the mortality of younger generations. I would qualify that. There was an especially interesting paper produced last year, published in the Journal of the American Medical Association, which found that in a very large study the mortality of people who were classed as overweight was actually lighter than the people who were classified as being of normal weight. This had actually caused people to re-evaluate what it means to be overweight and what the potential harm/benefit of carrying weight of a certain amount is. I am not arguing that people are who are very obese do not have excess mortality, but I think that there is a bit of a debate at the moment as to precisely how much effect increasing obesity and the increasing average weight of the population will have, especially given that heart disease mortality and stroke mortality are falling so rapidly at the same time.

My own personal view is I think factors such as statins and the effectiveness of blood pressure control in medication, and so on, will comfortably outweigh the negative impact of things like obesity.

THE CHAIRMAN: The other point that I would make in terms of interdependence is of course you also have the interdependence of medical research, and medical research will tend to be directed towards those areas where they can potentially have the greatest benefit and benefit in one sense related to mortality. Therefore, the areas where the medical research is directed, where we will get future improvements, may come in those areas that have the bigger impact as a result of other trends. It is a very, very complex model.

MR WILLETS: That is definitely true. That is one of the prime difficulties you have got when you are projecting mortality by cause of death, which people were suggesting this morning perhaps as a methodology.

It has lots of merits for different reasons. I think that it is very good for ICA work. It is also very good at picking up short-term projections. Where it starts to get really difficult is 20-30 years out we have reduced heart disease mortality massively; cancer mortality is also reduced by a large degree.

What then happens to the other causes of death? As Dave suggested, basically, medical resources will shift from where they are at the moment to these other causes of death. Modelling how that dynamic works is actually quite complex.

MR RICHARD ABRAHAMSON (Watson Wyatt): Richard, a fascinating talk with a rather alarming conclusion I thought. Thank you.

The question that I have is whether within your statistics, your work, you have been able to look at the influence of affluence, as somebody might have called it before, whether there is some sort of breakdown between different socio-economic groups which could be very important for application to pension funds particularly.

MR WILLETS: I think, generally speaking, over the past 30 or 40 years, if you explore trends in the English and Welsh population or UK population relative to CMI trends, for example, so people with pensions annuities, life assurance products, and so on, you definitely get a consistent pattern of faster improvements in the CMI experiences.

There are actually a couple of reasons why. They seem to be just a little bit faster across the age range. There also seems to be a slight difference in the cohort patterns for the two groups. You have slightly earlier improvements in the CMI trends showing that the healthy cohorts tend to be born earlier among the higher socio-economic class groups. That has caused part of the trend.

But, generally speaking, you have a consistent pattern of faster improvements for the high socio-economic classes. You can also see that in government statistics, which try to breakdown life expectancy and mortality rates into different socio-economic classes according to people's professions. You have exactly the same trend, more rapid improvements for the higher socio-economic class groups.

THE CHAIRMAN: One of Daniel Ryan’s graphs this morning showed the US and England and Wales starting at a very similar place roundabout 1960 and being in a similar place now, and had us lagging behind. I wonder to what extent that was because the US had perhaps picked up on the medical developments more quickly, whereas in the UK they have to go through an NHS approvals cycle. Therefore whether you can actually get different effects with changes in medical provision if people are more likely to be paying for their own care in future, they are more prepared to make those decisions for themselves than the government do collectively.

MR WILLETS: Perhaps also a factor as well is in the UK we have actually seen population trends and insurance trends being relatively similar in terms of their magnitude and shape. As I said just a minute ago, they are slightly faster in the CMI groups.

I think in the States there is more of a divergence between the average population rates of improvement and what you see for the more affluent groups. You have some of the highest life expectancies in the world for the more affluent groups in the US. But actually on a population level they are not particularly good in terms of mortality rates at different ages compared to other developed countries.

MS MELANIE CUSACK (Towers Perrin): A lot of this stuff is intuitive. Medical advancements make you live longer; there is less death from heart disease. But effectively in a lot of ways it is actually just delaying. If someone has a stroke they do not necessarily die at that point in time. But it could be a stroke-related death 20 years down the track, for example, or the heart disease is controlled through all the medical advancements.

Is there any point where it changes the classification, so people are deemed to have died from old age because they lasted until they were 90, but actually they were technically a heart disease victim and, years ago, they would have died?

I just wonder whether there is any evidence in the rates of improvement of that coming through.

MR WILLETS: There is a cause of death, I think it is called senility without psychosis. I think it is where they could not really find anything particularly wrong with the person, they were just very old. That particular cause of death is quite rare but it is becoming increasingly common. It is definitely a growth area in terms of cause of death!

MS CUSACK: Given that is a growing area in a sense, should we not be looking at that? It is saying medical improvements are happening -- it goes back to a comment earlier this morning – but physically, we have a finite lifespan. Okay, you can have all the medical improvements in the world but we will die at 120 -- maybe in ten years' time it is 130 or 1000.

MR WILLETS: It is certainly valid that if you look at trends in causes of death at high ages, say above 80, 85, and so on, we actually have quite a different dynamic to trends for people in their sixties, seventies and early eighties. There is much more ambiguity about the cause of death; the cause of death approach does not work particularly well at those very high ages. It is very difficult to tell in some cases what people have actually died from.

You almost need a completely different model to model mortality at these very high ages, sort of mid-ages, in terms of pensioners and annuitants.

THE CHAIRMAN: Thank you for the questions. May I ask you to join me in thanking Richard for his presentation? [Applause]

Taking Account of Uncertainty (1)

Communicating Uncertainty:

To a Life Company Board

THE CHAIRMAN: We now have another "doubleheader", if you like, on the topic of communicating uncertainty. John Lister is from Norwich Union and Nigel Bodie is from Watson Wyatt.

John Lister is currently Chief Actuary for the Norwich Union Life Group and a member of the Executive. He is responsible for the day to day running of the finance division and all actuarial matters, working closely with the Finance Director. John is a qualified actuary. He has held various roles and responsibilities within the Norwich Union Group and its predecessors over the past 18 years in both finance and marketing. In addition, he has chaired and been part of some Institute of Actuaries' and Industry working groups. He is a member of the Institute Council, sits on the Life Board and is involved in the Research and Education Committees. He is also an ex-colleague of mine, which is how I was able to 'lumber' him today! I will ask John to make his presentation.

MR JOHN LISTER: What Nigel and I have been asked to talk about is communicating to non-actuaries. I am conscious that the room is full of actuaries. So, hopefully, at the end of this you will have some idea as to how to talk to non-actuaries. In this respect, I am reminded that there is Peter Clark, who was President of the Institute between 2000 and 2002, who recently sadly died, had his theme for his Presidency of communication, culture and companionship. Communication was very much the dominant theme.

What I would say about communication is that it is not rocket science, but it is very hard to do effectively. The best tip I can leave you with is to think about your customer. Who are you presenting to? What do they want? What will they understand?

As a profession, we do not have a very good reputation for communicating effectively. There was a recent article in a newspaper where the editorial said "but he is an actuary not a communicator". I do not feel particularly good about that so hopefully after this session and Nigel's session we can all move forward. In the 10 to 15 minutes we have, I want to take you through some of the techniques that I have used and the thought processes that I use in communicating with the Norwich Union Life Board and the Aviva Board. If we do not communicate effectively it is our clients who will suffer.

At that point I am reminded of the actuary who is standing by a pool. He sees a man running towards him. He whispers to the man, "The pool is empty". The man dives in and of course hits the bottom and injures himself badly. The actuary shouts, "Told you so!"

What we do not want is our client saying "You never told me about the that." We want to be clear with our clients that they know where the risks are. It does not mean that we have all the answers, of course. Will come on to that in a minute.

I am also conscious, looking at my slides, that I have only two graphs and no formulae. That might give you a bit of a clue!

What I am going to talk about is starting with what is important to a Life Board? Nigel will cover what is important to the Trustees. Why would they be interested in longevity? I am going to talk about presenting results in an accessible format and, again thinking about the customers, where are the Board going to come from, because they are going to ask you some questions? At the end of our session, we are going to have some questions.

What do boards care about? If you think about your board what they are caring about is: am I maximising the value for my shareholders? Am I using my capital effectively? Am I preserving my reputation with my customers? That is about the long-term value of my company. Have I enough money to satisfy the regulators? If I have not, I have a problem with my customers.

The boards are accountable to their shareholders, their customers and their regulators. That can lead to a few conflicts of interest.

What you also need to think about with boards is that these are not actuaries but they are intelligent non-actuaries. They are not, as has been suggested in the past, a gentleman's club who is just there to pick up their money and are not accountable to anybody. That is very definitely not the case. They have some real priorities and they ask you some very tough questions. Anticipating those questions is quite important because there is nothing worse than standing there having to waffle, get into the theory too much, get into too much detail. My rule of thumb is KISS -- Keep It Simple Stupid -- me being the stupid one. They are definitely not stupid. That is how to lose your audience.

Preparation is the key. I am very conscious that this is a CMI seminar. Therefore I thought I would show you the graph you can now see on the screen and say that is not what they want. I am not suggesting that this is not right. One of the interesting things about mortality and longevity is that we do not know what is right, as Iain eloquently put it earlier. We know what the past is; we do not know what the future is.

What I am suggesting is very difficult for non-actuaries and some older actuaries like me to absorb in a few seconds what that means.

When communicating with boards, simplicity and clarity is the key. In my own organisation we use something called a SOAP, which is a Statement on a Page. Believe me, it is possible to get some quite complicated issues over in a statement on a page. I cannot remember whether it was Oscar Wilde or something else who said, "I wanted to write a short letter but I did not have the time." It takes time but, believe me, the communication is much more effective if you can get it on a page.

If you can get it on a page and it interests the board, they will read the bigger papers. If it does not interest them, you are lost. You will not grab their attention. I will swiftly move on, since I have probably insulted half the audience.

Why is longevity particularly difficult to communicate? In the session this morning Iain and his colleagues talked about how to come up with different methods of projecting mortality improvements. I do not need to add anything to that. It is incredibly complicated. Whether the P-spline or the Lee-Carter or the Weibull distribution, or anything else, is the right thing, I do not think I am qualified to make any judgments on.

But what I would say is what we need to do is to understand sufficiently and give it our best shot in the certainty that it will be wrong. So the key will be explaining to people the range of outcomes rather than the precision of the outcome.

Why should the board be interested apart from their interest in how they are going to live? One of the big three contributors to the ICA stress test is one of the reasons. Equity returns, credit stresses and mortality are three of the biggest stresses that we see, causes of risk across the industry. I have a £15 billion annuity portfolio to run. So understanding mortality, understanding credit and communicating that, and communicating the risks, is pretty critical to us. It is also one of those things that is quite difficult to do anything about.

I have put on the slide that you can see on the screen "little capital market activity". I think we have seen longevity bonds coming out but they have not really had a big impact in the industry yet. I have also referred to "lack of reinsurance capacity". That was certainly true, but the reassurers are beginning to get a bit more interested now. So maybe there are ways in which we can take some mitigating action.

But the key point is we do not know. Therefore what we need to do is write the correct volume of business to satisfy the directors from a "Have I got enough money to meet my solvency requirement, and how profitable is it?" point of view.

The fact that they are very long-term liabilities means that it is uncertain. We do not know. Richard has just talked about the impact of medical improvements. Ian was talking about the impact of pandemics, of changing economic circumstances, or whatever.

What we do know is that it is going to impact on the profitability of our business and the solvency of our business.

We do know that, in the past, we have consistently got it wrong. Hence what people believed was profitable is now unprofitable. So we need to be better but we cannot be certain. The challenge that you have with a board is that they remember the fact that you got it wrong in the past so why would they believe what you are telling them going forward? That is a problem of credibility and trust in the actuaries and in the profession as a whole.

The work that we are doing today in trying to estimate better than we did in the past is clearly important as part of the fact, as is the credibility of the CMI.

It is not that boards are not familiar with uncertainty. They understand the stock market volatility. They understand the fact that some companies fail. They understand that business expenses are volatile and that sales are volatile.

What they are concerned about with longevity is that they cannot do anything about it, apart from not write the business in the first place. If you think about that from a social policy position that is going to be quite difficult. As an insurance company, one of our USPs is taking risk yet we do not know how to manage that risk, therefore you should not right that risk, therefore do you have a business? You want to write this stuff but you have to accept that there is uncertainty.

One of the things that I try to do is to explain to the board and the senior management what impact it has on things that they understand. Impact on profit they pretty much understand, because it comes straight through in the underlying share price eventually. Impact on FSA reserves? That becomes really important when there is not enough capital in the fund and the shareholders have put money in. That is a good way of getting their attention.

The risk of ruin is the same point as before. Plausible future scenarios -- medical progress. Certainly one of tools that we have used in explaining things to the board is “This level of improvement is consistent with a cure for heart disease or a cure for lung cancer”. So, trying to relate it back to something that they would understand, as is trying to give things like expectations of life. It is a very simple way of showing how people's expectations of life have improved over the years to give them a reference point to think about what is likely to happen in the future.

As I have said, none of this is rocket science. I should point out that the numbers you can see on the screen at the bottom of the slide in terms of expectation of life assume no future mortality improvement. If you assume future mortality improvement you would get a slightly different picture.

Richard has covered the medical advances. There are also things like pandemics. That is a good one. Boards are very conscious of what is going on in the world. In fact, they would be very good at answering questions on the News Quiz, or something like that, because they do tend to read a lot of analysts' reports and newspapers. I will come onto that in a minute. When you talk to them about

ICA and the potential for mortality improvement to carry on for our long period of time, they will say, "Ah! What about avian 'flu? That is going to be good for us, is it not?" Or "What about all these kids who eat too many McDonalds and do not exercise? What about global warming? Is that not going to impact on life expectancy?"

All of those things are perfectly sensible. They are challenges that you will get from boards when you discuss it with them.

One of my colleagues who is in the audience picked out the following. This is from the BBC website, some titles or headlines. As you can see, these things are going to influence people. I particularly liked the one that you can see in the middle of the slide which says "Yeast intake linked to longevity". Anybody who deals with the press will know that the headline has nothing to do with the article. This article is about the impact of yeast on fruit flies, but you would not know that!

What I am trying to do here is to give you an idea of what forms boards' opinions.

They also read the analysts' reports. I particularly liked the one which said "insurers assessment of their our [annuity] liabilities is likely to be realistic". They obviously have not been listening to Richard talking about minimum improvement rates. I was quite gratified by the Bernstein[?] report; but, if I want to strengthen my mortality to reflect some of Richard’s views, that is going to cause me a bit of a problem. The analysts are saying it is okay. Why am I saying it is not okay? If you look at this there is a range of views which causes confusion. You are there to try to help with the confusion or to clarify it.

I talked earlier, do not give point estimates. The boards are comfortable with the concept of a range of outcomes. You can make statements like, "We are 90% confident that it is in the range of £19 billion to £21 billion" but you are likely to get the question, "Well, can't I use £19 billion, then, rather than £21 billion?" You need to start getting into the different uses. Richard alluded to this. You need to start talking to them about what is the purpose that you are using this for. Is it profit? Is it solvency? Is it your ICA calculations? Try to get across the concept that you need different levels of prudence at different levels.

The best piece of advice I can give you, other than think about your customers, is "say it with pictures". People can understand graphical representation very quickly and absorb it and understand it.

The picture you can see on the screen gives you a good idea as to the range of outcomes that we might see over the years from 2% mortality improvement to 3½%. That helps them try to understand where they should be when you are using this for different purposes. Wherever possible, use pictures.

I have to say, based on my experience, that the board will turn off at anything technical. If we started to try to explain to them the different statistical techniques for analysing mortality and forecasting mortality, I am afraid that they would lose interest very quickly apart from the one person who did a degree in this 30 years ago, who is going to ask you a lot of difficult questions. You get into a one-on-one debate that does not really work for the rest of the team. I would not get into that, if I were you.

You need to accept that other results are possible while at the same time not undermining the message. You can get into a position if you say "There is a range of outcomes and I do not really know where it is going to be, etc" that they lose trust and confidence in you. One of the observations that I would make is that it is pretty difficult to come up with something that we can all agree on in terms of forecasting longevity because we do not know.

But not undermining the message is quite important. You have to be prepared to answer the question, "What do you think, then?" because you are giving a large range, and they are paying you a lot of money to give them advice.

So it is a tricky balance to make between saying, "There is a range, but I would tend to go up here, and these are the reasons that I would tend to go up here". It is very important to maintain their confidence.

It is also very important to be professionally happy with the decisions that are reached. There is nothing worse than coming away from talking to a board unhappy with the decision that the board have made for you personally and from the company and the policyholders' perspective. So you do get into conflicts between policyholders’ and shareholders' interests, and you need to be very clear professionally where you are going to draw line and where there is an acceptable range of outcomes.

So, how would I summarise all this? I would say we can get the message across but we have to address their priorities. If they are not interested, you will not get anywhere. Make it clear and concise. If you can get it on a page, do so. If you can use pictures, do so. If you can simplify, do so. Anticipate their concerns and challenges.

Remember that they read newspapers; they read the analysts' reports. They have concerns about sovereignty of the fund and the profitability of the fund. Make sure that they trust you and that you can demonstrate that trust and credibility and the advice that you are giving.

Before concluding, I should just like to quote from Peter Clark's presidential address. Peter finished with: “The challenge is to communicate, not in mathematical symbolism but in persuasive words. Let us ensure ‘I am an actuary and I should like to explain’ an accepted reality not one of those four unbelievable sayings."

At this point, I should like to hand over to Nigel to cover the angle from a pension trustee’s perspective. We will come back to questions at the end.

THE CHAIRMAN: Nigel Bodie is a senior consultant with Watson Wyatt and holds a number of scheme actuary appointments as well as having a research interest in longevity and valuation methods and assumptions. He is a member of the Current Issues Committee of the Pensions Board and has recently retired from the Council of the Pensions Management Institute, where he was a Vice-President. Nigel is also a member of the CMI Self-Administered Pension Schemes Committee of which we heard earlier.

Communicating Uncertainty:

To Pension Schemes and Trustees

MR NIGEL BODIE: Just to start off with a nugget of information for which I am grateful to the BBC. Today is Siesta Awareness Day. I did not know this. If I am not on their mailing list, who is! I will see what I can do to promote it during the next 20 minutes.

We are talking to a group of people who know that there is an issue here, mortality and improvements; they are keen to understand what is going on. It is a major item of uncertainty for them; and for ordinary companies, not life offices, but for companies it has not really struck them before and they really want to know about it and they want to understand a bit about what is going on.

The actuary's problem falls into two parts. One is getting current mortality right, to which the comment is, "Surely, that's easy." The second one is "What about future improvements, and why can't you actuaries get it right?"

I am going to come up with a few ideas as to how you might communicate issues and answers to these questions. I will go through it quite quickly. In an audience such as this I do not need to explain and what an A/E is, but when we get to talking to boards of trustees then certainly you need to take more time to explain the issues.

Let us have a look first, then, at taking a snapshot. The opening snapshot is easy. It is actually drawing the information out of it and using it to make future projections that is the difficult part. You can start of in many cases by saying maybe the current picture is not even representative. If you are looking at a large company which has changed the nature of its work, you may find that you have a pensioner population which is largely blue-collar; but for the population that is coming forward for whom you are valuing the actives, you have a completely different white-collar population, and so your existing population is no help in projecting what is going to happen to the others.

There are other sources of uncertainty. We have had the low and high paid; industrial variations; and regional variations. I will touch on all of those in a moment, just giving examples of how one might illustrate this.

To start with, you have even got year-on-year volatility. I have displayed on the screen a chart of just the CMI's report on its own investigation, male normal retirements compared against 92 series based table, and even if we knew where the trend line was - with the benefit of hindsight of course we do - you can still see that from year to year there are significant variations in the rates of observed mortality.

This comes as a surprise to trustees. They have some vague feeling that there may be good and bad winters, but the extent of the variation, literally from year to year, comes as something of a surprise. So you can start off with that at a global level. This of course is even when you have got a very large amount of data.

Large and small schemes alike will tend to say "Can't you just use the standard tables?" This is quite a popular question to which the answer is, obviously, "No". We can demonstrate why not because we then look at how the standard tables compare with the experience of the large occupational pension schemes, the schemes that Brian was talking about this morning.

First of all, we start off with a comparison of the data for male normal retirements against the 00 series. Even though we can get the A/E down at 99%, pretty much spot on, clearly the shape is not very impressive, particularly at the younger pensioner ages. Even here we are having to use a one-year age rating, so for the 00 series -- this data is centred around the 2001-2002 -- we are having to use a plus age rating on it even though we are in fact using the wrong table, so to speak. This table is not a particularly good fit. This standard table does not work for schemes in this particular group as a whole.

Similarly, the 92 series table has much the same thing, except that that the tweak is at that the other end. Again, this is against the base table. This is 92 mortality rated minus a half. So we get a A/E of 99% again but the shape is completely wrong. If you push in the cohort improvements to get the factors around age 70 right, then you push the younger ages even further out of step. So this is again not a very good fit.

There will be some of you out there who will be thinking "Why hasn't he chosen the obvious table, that is 80 series calendar year 2020 rated minus a half?" And I applaud your perspicacity! You are absolutely right, the old ways are sometimes the best.

A word of warning at this point. Please be careful when using this sort of illustration in the presence of pregnant women! The lady in charge of our transfer value team called on me one day. She said "We are reviewing the default basis that we use for transfer calculations. Can we use a default mortality table?" Bearing in mind this has to be published, this is set out as part of the scrutiny of the transfer value basis, I was able to say "Well, quite frankly, the two latest series are completely useless, but if you use this old table projected to an arbitrary yield and an apparently arbitrary age rating, it is absolutely spot on and I am sure that you will have no difficulty explaining that." That was half past six that evening. Later that evening her baby was clearly overwhelmed with excitement at this and could not wait to put in an appearance and came out six weeks early. So please be very careful when pregnant women are around you; you do not know what you might be doing!

Variation by amount: Brian touched on this also this morning. You can see on the screen the amounts data. We have a little bit of a problem on the SAPS Committee looking at the data at the younger pensioner ages. There are certain inconsistencies which we are not quite sure about. I have adopted the age old actuarial technique of ignoring anything that does not support the story that I am telling and just showing the data from 65 onwards. Again, you can stun the Trustees by the extent of the differential in the mortality rates between the low and high paid. Around age 75 it is double. They cannot believe it is as significant is that. It is also interesting just a note how the lines come together at the older ages which begin to imply that there is perhaps something more to do with luck and genetics once you get to these ages, but at the earlier ages there is a significant amounts effect, and they are genuinely surprised at the extent of it.

So you are beginning to get them on your side now, they are beginning to see that maybe this is not quite as easy as we thought.

Another popular thing is regional variations. For a particular client, a large client which had offices all over the place in the country, we investigated their mortality experience. We looked at first of all the raw mortality rate. Clearly, there are the expected bad hotspots of mortality in Glasgow, Birmingham, Eastbourne and Worthing! Once you allow for age, year of birth, sex and amount of pension, there is not a lot left. There is still some effect from geographical locations but a lot of it is actually described by some of the other variables.

You will notice that with all this data and all the smoothing that is going on there is still one little spot down in

Penzance which really appears to be the place to go if you want to live a long time.

Also in the SAPS investigation we looked at variation by industry. We have listed on the slide that you can see five of the industries that we had. I have ranked them in the size of the average pension. If the average pension is rising, you would expect the A/E to be falling because you are looking at higher paid people for whom we were expecting lower mortality rates.

While there is a general trend for it to be from high to low, it is by no means consistent and there are significant industry variations which make again the adoption of the standard table completely inappropriate.

You really have the trustees on your side now. They say even taking the snapshot is not as easy as they thought it was going to be. So now the question is "Where are we going?" Instead of asking the question rather aggressively, they are now rather interested -- "Where are we going now? What is happening?" We can point out that there are variations between experts; we can point out the cohort effect; and we can point out trend uncertainty.

Some of these matters have been mentioned before so I will not go into them in too much detail. James Vaupel thinks that the mortality is going to keep on improving as we go. He plotted since 1840 the country that has the greatest longevity and then put that on the graph that you can see on the screen showing the longest life expectancy. It forms a staggeringly straight line. He draws in the projections that people made as to where the limit of life would be. He is particularly proud of this one which was published and it was wrong even at the time that it was published. A number of the others got overtaken as they went on, but he was wrong before he started! Vaupel said that this is just going to keep on going.

Olshanksy, on the other hand, we have mentioned, does believe that in some sense there are biomechanical limits on lifespan. He accepts that we might be coming more rectangular in the survivorship curve. The extent to which we are pushing out the bottom end may be a little more limited. His argument is that there will be more people surviving into their nineties but the limit to life may not be moving.

With regard to other experts, Professor Tom Kirkwood gave a talk at the NAPF recently. He said that there is no strict biological programme for ageing and no limit to the length of human life. He again made the point that we are likely to see more people pushing through into the nineties before we see too much improvement right at the full expectation of the limit of life.

Dr Aubrey de Grey, according to the BBC website, said "I think the first person to live to 1000 might be 60 already". That is an even more sweeping statement than "has just been born". Dr Aubrey De Grey is a very good publicist. He needs to raise money for his research work but he is very interesting and a fascinating guy to listen to.

The cohort effect. This is where you really can hit them because trustees are very surprised at the extent and the duration of the cohort effect. There are a number of figures in Richard's paper, for example, which show how this cohort marches through the population, with this particular group showing better improvements than anybody else. You can illustrate how that has happened in the past and they would again be genuinely stunned at the existence of this cohort and the significance of it, how much more it is doing much more than other groups.

If you are going to go down this route, make sure you have read all Richard's papers because the next question that there will ask you is why is this happening? This is quite difficult, but there are some ideas that Richard has put out which will give some answers that you can give.

What is the effect of it going forwards? You need to give some sort of graphical illustrations. John mentioned that people will understand the chart much better.

This slide to shows some Qx tables at age 80 and 90. 80 is quite a nice example. You will see why. It is worth commenting on the uncertainty as to how long this will go on. We have not a complete explanation of the cohort effect so we do not know how long it is going to happen.

Looking at a chart which might give some indication of how you might illustrate the Qx, if we look at age 80 the base table improvements are shown on the green line that you can see. A nice steady decline in Qx. All the other three cohorts are there and you can identify them. They run fairly much together at some periods but each one ends up in a different place. It is worth showing that what is happening here is that this group is coming through. They are the favoured cohort. They are pushing the Qx down and we do not assume that they go back up again to the original trend line. This bunch have pushed the Qx down and everybody coming after them benefits from that. This chart shows that very clearly. It is just the extent to which it happens. People quite often do not understand. They think that this particular group are going to have special annuity values but everyone else is unchanged. That is not the case. This group are affecting everybody who comes after them as well.

Finally, if you have a really expert bunch of trustees you can lift this chart out of CMIR 15. This is a sample annuity portfolio and the reserving requirements for it. This is using, I understand, a P-spline model to put limits on the size of reserves that you should be holding...

[TAPE CHANGED AT THIS POINT]

...100, the 95^{th} percentile limits are 92 up to about 112, just on Stage A, which is the part that shows only variation about the trend. You do not know where the trend is going, so where there is trend uncertainty the reserves could be between 92 and 112.

The other stages add various levels of uncertainty such as variation around the trend line, if you knew what the trend line was, variation between people on high and low pensions. Then, finally, showing that people on high and low pensions have different mortality rates as well. All that is getting a bit advanced; but, nonetheless, with time and with people really interested in understanding what is going on, and the effect of that uncertainty on reserving requirements, they will come with you, and also pointing out of course where you have relatively large numbers of people, you get a group of about 2500 pensioners, then all the other uncertainties go away and it is only the direction of improvement which is the thing that causes the real difficulty.

That is just a canter through just a few points that may be helpful in getting the trustees back on your side when they really thought that the actuarial profession did not know what it was talking about on mortality.

THE CHAIRMAN: Thank you very much Nigel, and thank you John again. Some very helpful views, there. I detected a slight divergence. John, I felt you would not have included the last graph that we saw on the screen in the presentation to a life board.

MR LISTER: I do not think I would.

THE CHAIRMAN: And one of those tips was too late for me because I was speaking to a company about its pension scheme last week and, yes, the Finance Director did have a doctorate in medical statistics!

Questions from the audience, please?

[?]: Thanks, Nigel, for your presentation. I am interested in pension schemes, so it was very topical. I still feel that there is a gap. We have the client onside, they believe that we are not hiding anything from them but we still have quite a leap to make for the trustees to come up with a mortality assumption. They understand the difficulty now, but they still look to us to say, "Okay, that is fine. We understand that it is tough, but what is the answer?" I think we still have a long way to go to be able to say, "And, given that you are a small scheme" as you have said without experience, or whatever, "this is now where we need to go." I think that we as a profession cannot actually say any more: "This is what we would do." They now have to say "This is what we are going to do with the actuarial advice is and education that we have needed." Any comments on that?

MR BODIE: In the end, now with the Pensions Act, with the trustees ultimately responsible for all the assumptions that are being made, I think if you have meaningful data then we do not have much excuse for getting the starting point wrong. We should be able to make a pretty good shot at that. Thereafter, I think that we are being asked to say what is a reasonable provision for future improvements? We do not know what it is; we do not know what the reality will be. We can only do something that is reasonable. The regulator has specifically mentioned the provision for future mortality improvements. I suspect it is because he has seen or heard of a number of schemes where, frankly, they have more or less ignored mortality improvements even now, so they are pressing to make sure that people are doing something.

Again, I think that we have a problem in that, for example, in accounting figures I have been asked from time to time, or colleagues have been asked, by the auditors, "Is this your best estimate of future improvements?" You may find it hard to believe that I dissemble in answering this one and say "I have no better estimate than this particular one" whichever one it is that I happened to have chosen! As one of my more technical colleagues said, "What you mean is you have a uniform probability distribution".

I think that we are in a position of saying do something reasonable. No one is going to criticise you for getting it wrong so long has you did something that you thought was plausible.

THE CHAIRMAN: The question related to pensions, but the board increasingly takes decisions in a life company.

MR LISTER: Yes, the position is not any different in reality in that the boards of life companies have to take responsibility for the assumptions. You are absolutely right, we have to come off the fence and to say "This is what we think we should be doing. This is where we think we should position future rates for mortality improvement." Just like Nigel, we have challenges from the auditors as to what they think might be the right one. In the case of the ICA, we have the FSA challenging us as well.

We are being paid to help boards make decisions. As I tried to draw out, the level of prudence comes down to what is the purpose? If the purpose is to come up with the best estimate, then give them your best estimate. If the purpose is to come up with something prudential, then give them some idea of where you might put the prudential rate of mortality improvement.

On that point, the last point that Richard left us with was a minimum rate of improvement. I agree entirely with what Richard said. I think if we do not have a minimum in there, we are really not in the right place.

MR MIKE HARRISON (Mercer): I can understand the CMI does not want to give recommendations or call anything standard tables, but at the end of the day, we do need to recommend things when we are actually dealing with clients. FRS 17 is an interesting example. I liked Mr Willets's comment at the end of his presentation saying that the MC tables were probably too weak.

My experience is that there is an awful lot of client pressure to have something much, much weaker than that to give clients the answer that they want, and they do not even have to disclose what the mortality assumption is.

There is a new version of FRS 17 which has been drafted which says there should be some disclosure regarding the mortality assumption. To what extent is the CMI and all the profession working with the accounting profession to determine how that disclosure should be written to ensure there is sensible comparability between companies? If we are not working with the accounting profession, why not?

[?]: The short answer is the CMI is not, and I am not sure that the CMI would be the right arm of the profession to take the lead in negotiations with the accountancy bodies. It is a good question as to which part of the profession should be taking the lead in that.

THE CHAIRMAN: I think that the Life Board and the Pensions Board and other boards are trying to present a more joined up picture to different professions. Perhaps Ian would like to answer.

MR IAN MORAN: I chair the profession's accounting committee that cuts across all the boards. On that point, are we doing something? Yes. What is the answer? The answer is that we are promoting greater disclosure. The arguments for not disclosing such an important assumption would actually not do the profession any good at all. That is the view that we have taken. That is not just the life people on the Committee or the general insurance people, it is the pensions people and the life and the general insurance people.

So, yes, it is being addressed as part of that consultation paper. That is primarily in response to the needs and arguments set out by investors and other stakeholders.

THE CHAIRMAN: If there are no further questions, please thank Richard, John and Nigel for their presentations. [Applause]

Taking Account of Uncertainty (2):

Regulators View

THE CHAIRMAN: For the final session of the afternoon we are starting with Nigel Bankhead. We had originally hoped to "twin" Nigel with a presenter from either the FSA or the pensions regulator. Unfortunately, neither obliged.

The FSA view, which I received in an e-mail, so I can quote it, was "We are unwilling to put forward a speaker, as part of our ICA philosophy is that each firm analyses its own risks and then considers what stresses are appropriate, and we do not want to appear to be setting benchmarks for mortality or other risks." I think that is quite understandable. It clearly tells us that it is our problem and not theirs.

Nigel Bankhead qualified as a Fellow of the Institute of Actuaries in 1983. He was previously head of actual practice for Aon. He has extensive commercial experience and has been active in the development of professional standards. He was previously a main committee member of the Association of Consulting Actuaries. He is now Director of the Board for Actuarial Standards.

MR NIGEL BANKHEAD: When I received this invitation to speak I was tempted to do exactly what the FSA did, but as a standard setter you cannot really dodge invitations to talk about standards or issues that might affect them. I should like to say a few words about the FRC and BAS.

Turning to the FRC, I am sure that, as most of you will be aware, the FRC is the UK's unified regulator for financial reporting. What it is responsible for is everything covering accounting and actuarial standards. Its philosophy is based on the premise you can see on the slide on the screen: a well-informed market is the best regulator. Broadly, that translates into what you should do is simply provide information fairly, openly, comprehensively and objectively to the market and let it make its own mind up. So we are not here to set value; we are actually here to provide information and to set standards for that.

In terms of the debate which I am sure you have seen, over Enron and other things, and the principles versus rules, where we stand is that we actually believe in both but we do believe very strongly in principles. Going with that, that means that what we believe in are things like professional judgment. In actual fact there is a strong role for a practitioner to be exercising judgment based on some agreed principles. It is not simply everybody do exactly what they like, but there are some inherent principles which underlie what we all collectively would do. That is the basis behind it.

That may get translated into rules in some areas. Most things that we see are probably a combination of both.

Turning now to BAS, the first thing to say is that it is very early days for us. If someone came an asked me where are we on mortality, “we have not got there yet” is the straight answer. What we are doing is starting out on a development programme. The first stage of that will be establishing a conceptual framework. What are the basic generic principles which underlie actuarial practice that we all have in common and pretty much ought to bind us together as a profession?

To give you one example, what might underlie, as a philosophy, the measurement of value? How do we do it? If you take it from accountants they used to have historic cost, broadly, and then they have moved to what they call fair value or market value. Is that our philosophy as Actuaries? Are we trying to reproduce market value and, where there is not a market, simulate a market value? Do we have alternative philosophies that we can develop, expand and then, the necessary part, persuade everybody else? It cannot be a philosophy that lives in actuarial land. It would have to be one that we could persuade the outside financial world that here is a philosophy that they should adopt.

Consideration of mortality fits into the category of assumption. That is going to come later. Generic principles first.

However, rather than end my talk here, which is very tempting to do, and to run away, what I thought I would do is to share with you a few initial thoughts that I have had which relate to mortality. The first thing to say is that they are not all about mortality. At the very beginning we have what is actuarial practice? What is its purpose? What it seems to be, thinking about it, is providing information to assist planning for a variable future. That is pretty much the activity that we are in. What that future seems to involve in terms of the areas where we work are generally financial payments. Most institutions we work with provide contractual payments. They are paid based on some contingencies and we have thrown in a bit of discounting. I hope that I have not missed out too much in that. I know that I am being fairly general.

What it does not involve though is also as important in my mind. It is not about fortune telling. That is going to be a really big challenge. I am sure the outside world think that we are fortune tellers, and what we are sure about is that is not what we are doing. There is a big distinction between providing sensible estimates for the future and actually trying to be in the business of fortune telling.

The other thing that it is not about is making the future less variable. No prediction that you make generally makes the future and what happens any less uncertain.

So, just a quick run through: financial payments and contingencies. As we all know, most of these payments affect people. They are things like premiums, annuities, lump sums, sickness benefits, the whole array, medical expense, and most are affected by death. When death happens, either the person gets paid or he ceases to pay you. It is pretty simple. So in most of the work that actuaries do concerning people, apart from things like driving cars, mortality issues are bound to be of high importance. So in the area of life and pensions mortality will be a very key area.

I have taken a couple of concepts – again, slightly related to mortality. I looked at this general business of insurance and I have thrown in quite generically pensions and said what principle do we get from that? It is broadly a pooling concept. That is the basis on which it is set up. You as a participant can exchange individual risk, or the impact of it, for a share of group risk. I have put impact in just for a laugh, just to make it clear. If you buy an insurance policy it does not stop you dying. I am sure we all know that.

Going on to actuarial practice, very broadly what we have is a very straightforward principle of extrapolation. That works on the basis that past experience provides a guide to the future. It is pretty much that simple. Most of what I heard this morning was really about that question, "what sort of guide is it to the future? What future will it be?"

On the next slide I went through and posed a couple of questions which I answered for myself. They were true or false and I put a tick if they were true. I will run through them and see is actuaries what we were doing and what we might have done. Mortality concerns risk. I think that gets a tick. Mortality is uncertain. Risk concerns events that are variable. That also gets a tick. Greater understanding is gained by studying past experience. That gets a tick. We do that so that is good. Greater reliability is achieved by increasing the data studied. The more data you studied, the better and the more reliable that information is probably going to be.

Here is an important point: the average outcome from a group will be closer to the true average; take a small group in the overall population, it is going to be closer to the national average as group size increases. That is quite an important concept. If you turn it round, what it actually tells you is the experience of a group is likely to be further from the average as the group size decreases. A point to note. We will come back and see that later.

Actuaries have treated mortality as fixed probabilities. I think that is true. Although it is very variable, we just get a set of tables that gives us a point number at each edge and we read it off. The result of that is, when you translate it, we have treated variables as certainties in the area of mortality.

You may agree or disagree, and these are not intended to be final conclusions at all.

Some challenges I see coming out of this area. The first one, which will be quite interesting, is how the past should be relied on when we come round to look at it. Is it averages? We say in the past average repeats itself. We say in the past trend repeats itself. We say in the past dispersion repeats itself. We are going to need to answer that question. If we are projecting mortality that is to do with trends. In some areas you may say, assuming a trend continues is right, that is sound, and there may be other areas where it is not sound. Should a rise in P/E ratio continuously repeat itself if you are looking at the price of equities? That is a trend and it is probably a very different trend.

Another question which is of interest to me: when does pooling break down? It breaks down if you do not have a pool. If you have set up an insurance company, a mutual, and you are the only participant, you have really achieved nothing in that insurance concept. But is there are certain minimum size that you need legitimately to run a pool? How does dispersion increase as that pool size reduces?

I guess a question going with that is when we are advising people with small pool sizes are we giving them information on that? That is going to be an issue.

We then come to best estimates. We talk about best estimates; it can be best estimate for anything. All that means is that you are trying not to cheat, so far as I can see, when you are giving a communication. We somehow have associated that with averages. Best estimate to us implicitly means average. It could mean anything. Why is that? It is limitations on computational ability, analysis, or the ability of clients to comprehend? I do not know, but it is the sort of issue that I think we need to look at.

Finally here I have posed the question: is actuarial practice concerned with individuals or pools? We tend to look at mortality on an individual basis, often. Actually, we are not in that business. Very few of us advise individuals on mortality. We actually all advise pools. That is the business that we are in.

I would ask myself a couple of questions, and there are no particular conclusions to these. But how does mortality affect the experience of pools? It is obviously to do with whom you let in, the entry conditions, perhaps exit conditions when you kick them out; what happens to their mortality, particularly pool size; also things like hedging, whether you can offset the risk or you cannot; things like risk exclusions. That will be something that we need to look at.

When we come round to the information, the main concern for us for BAS is actuarial practice, as we said at the start, is information given to somebody who uses that. They do not want to consume it in its own right. It is used as a means to the end to help them run the entities for which they are responsible. We will need to pose a question: what information is it that they need? Do they need information on the expected level or trend; the expected dispersion? Do they need information on disaster scenarios? Not saying that the average is here and it may be here but how bad could it be? Perhaps something like the FSA 1 in 200 year event.

Also, probably they need information on the impact of potential management actions. What action can you take if you are running mortality risk and what impact will that have on the experience of the pool?

That leaves me with probably the three questions that I ended up with, things to think about, perhaps a reiteration of what I have said already. What information should be provided by actuaries to people running pools which depend on mortality risk? How should the information be calculated; and how should it be communicated?

I think that that is pretty much it in terms of at the big picture level. As I said, those are my initial thoughts. At the moment we do not have the answers. We do not even have a full set of the questions that we are going to ask. That was very much just a few introductory thoughts into some of the sort of areas into which we should make some inquiries.

Questions

THE CHAIRMAN: How do you see the existence of BAS affecting practice in this area until such time as you do look at mortality? It almost feels like there could be a vacuum there.

MR BANKHEAD: That is a good question. How would I see the existence of BAS affecting that practice? At the moment, the first thing that BAS did was to adopt all the existing guidance notes which were in place related to what was regarded to be technical areas as opposed to ethical ones. So in any technical areas actuaries should continue to operate based on the guidance that is in place. I do not see there being a vacuum. You are right there is no change. The previous position, everything you previously did and the constraints you were subject to, you are still subject to. It is not as if the guidance was cancelled. It simply remains in place.

You may say, as I might well say myself, "Show me the piece of guidance in a guidance note which tells me which mortality table to use" or "Show me the piece of guidance in a guidance note which tells me how to determine when I am given choice". We believe in that professional judgment. But what is the process an actuary should run through to determine the right mortality table to use and where is that set down and documented?

To my knowledge, neither of those are there at the moment. It ought to be there. I think it could be criticised because it is not there. On day one it is no different than it was previously.

[?]: (inaudible)

MR BANKHEAD: There are some transitional arrangements. So, for example, when we started, the profession was running through a process of revising certain guidance notes. What we agreed was in actual fact that those guidance notes were at a very final stage of revision. They would finish and we would consider them for adoption when that was completed. If any new issues come up which are urgent issues, we would have to deal with it.

Fortunately, I think a lot of matters in respect of pensions have just been through a major process of change and the one that is going on there is largely transfer values, which is just starting out in terms of that debate. I think also in life insurance there has also been quite extensive revisions. You are quite right, if something came up that critically required to be addressed immediately, we would probably have to do that. The one thing that I would say is what we do want to do is to make sure that we have consistency between the different areas of actuarial practice.

Previously, there were three or four separate publishing bodies. Each board produced standards. You may not have been able necessarily to reconcile them. We do want have common principles so that where there are standards in life insurance, pensions or general insurance, it would, in actual fact, follow some common core principles.

[?]: [inaudible]

MR BANKHEAD: I am not saying it is an easy run in. But I think in terms priorities the priorities are to start with some of the fundamentals rather than some of the detail. Even if there are priorities, that would be the minimum fix that we have to do. We ought to start with the core issues first. It would be silly to start off with individual subject area matters and when you put the whole thing together you have a set of standards which are inconsistent, are not coherent. That is generally not helpful.

I take your point and I am aware that in life insurance there may be more specific guidance. The role of the FSA is different. You could argue that they have almost become a standard setter. I would not say that they like that, from the conversations that we have had with them. They would be reasonably interested at some stage in the future whether BAS could act to replace some of the standard setting that they have done. That still remains to be seen. That will be a process of evolution between different standard setters, the FSA primarily the prudential standard setter, us, very much a standard setter for information and calculations. Thank you, and I will look for that life guidance.

THE CHAIRMAN: We will move on now to the final presentation.

Taking Account of Uncertainty (2):

Modelling and Measuring Uncertainty

THE CHAIRMAN: Steven Haberman is Professor of Actuarial Science and Deputy Dean of Cass Business School, City University. He has previously worked at the Prudential Assurance and for the Government Actuary's Department. He has been a member of the Council of the Institute of Actuaries. He has also recently been appointed to the FRC’s board of actuarial standards and to the ABI's research advisory panel. He was a member of the external advisory panel to the Morris Review. He is co-author of four books and over 140 papers. I do not intend listing them. They have won research prizes from the Institute and also from the Society of Actuaries in the US.

Professor Richard Verrall is Head of the Faculty of Actuarial Science and Statistics at City University and is Associate Dean at the Cass Business School. He is also the author of an impressive array of papers, many of which relate to claim reserving and general insurance, but also encompassing mortality projection models.

I will hand over to Steven first.

PROFESSOR STEVEN HABERMAN: It is clear that you have had an interesting and a long day. Please bear with us, I think that we are going to be discussing an interesting subject; namely, the modelling and measuring of uncertainty. We are going to do this in the context of the Lee-Carter model simply because we need some bench mark to use for showing you how to deal with uncertainty. The techniques which Richard will be talking about are general, generic techniques that could be used in any of many contexts.

I am going to start off by talking a little bit about Lee-Carter and some of the recent developments. I know that you have already had a presentation on Lee-Carter. Please bear with me, there will be a little bit of duplication but I hope that there will be one or two things that you may not have heard this morning.

Some things about the Lee-Carter structure: globally this has become something of a standard. A number of countries now are using the Lee-Carter methodology for modelling and then forecasting the trends in mortality rates, particularly the US Bureau of Census uses it as its benchmark. No doubt, as you heard this morning, the structure is deceptively simple but then turns out to be very complicated. The simplicity is that we propose that the mortality rate, the log of the mortality rate, is made up of two bits, something that we try to predict and then an error term, and the predictor itself is made up of some simple components: an age effect, which has come to be known as alpha; and then a product of terms: kappa, which represents time trend, and beta, which represents the age response to that time trend. This is a representation of the mortality rate at age x and year t, decomposed in this way.

It looks simple. The complexity comes because when we immediately think of regression we have an observable quantity on the left but no observable quantities on the right, which is rather unusual. There are some constraints that people impose in order to deal with some indeterminacy in the model.

When Lee-Carter first proposed this in 1992 -- Ronald Lee and Lawrence Carter -- they used a way of fitting called singular value decomposition, which is a bit unusual. If you are familiar with multivariate analysis, it is a little bit like principal component analysis.

Since then more sophisticated and more modern techniques have been developed based on weighted least squares using an iterative method and using the more conventional maximum likelihood methodology with Poisson errors, as we are used to using in graduation.

There is an adjustment made to the kappa values once the fitting has been done so that the graduation property of actual and expected deaths in each year being equal is satisfied.

Also, what has happened over time, as people have used Lee-Carter over the past 13 years, they have started to impose the sort of things that we do naturally as actuaries to measure goodness of fit. Looking at residuals, which was not something that Lee-Carter first proposed, they had some very simplistic ways of deciding whether it was a good or bad fit, and looking at patterns of residuals is a very powerful methodology. Because of that when Lee-Carter is being used in certain countries, in particular England and Wales, it is found not to fit the data very well.

It turns out to fit pretty well female data but not so well male data. So some enhancements have been introduced over recent years. One is to add extra terms, an obvious thing to try to do to the model. The second is to reflect on what we know from the mortality experience in the UK;

namely, that there is a cohort effect. It is not true of all countries but it is certainly true in the UK.

I will show you how that is done because I think that is quite dramatic. Forecasting is done, as I am sure you heard this morning, by fitting a time series model to this time trend. We get best estimates of the alphas the betas and the kappas and then we just focus on the kappas and try to fit time series models using the full panoply of ARIMA models or, if you like, econometric modelling. There are many packages that enable you to do that.

Then you can get a forecasted estimate for future mortality rates. You can see on the slide on the screen and at age x year t n is the last observed year but moving s steps forward.

A cohort version: given that we know there is a cohort effect, how might one construct a cohort version of Lee-Carter? One could say you have the alpha plus beta times kappa, why do we not add a term that represents year of birth? So we have alpha plus a beta times iota (iota is a cohort term) plus another beta times kappa where kappa is a calendar year term.

If one could fit that, one would have two things to try to project: model the iotas using time series methods; model the kappas using time series methods. There is an immediate problem as most people notice when they think of cohort methods. That is you have your simple relationship between and cohort period and age. You know the year that you are looking at, you know the age, and you immediately get the cohort just by subtraction. That requires a two-stage fitting strategy.

What we have tried, and it seems to work well, is to estimate the alphas, which have an overall averaging effect. What does the overall age profile look like across all years? Fairly stable. We estimate that first, get that out of the way, and then estimate the betas, iotas and kappas.

It is quite demanding in terms of data so we have used a data set provided by GAD, which covers a very long period of time, over 50 years, and provides data by single year of age. We have also done it using CMI data for assured lives. Similarly, a long data set enabling us also to use single year of age.

We have not been able to use it for annuitants because the data set is too short and does not give us enough data in order to identify cohort effects.

The purpose of the mass of dots that you can see on the screen is to show you what residual patterns look like. This is obtained from just fitting the Lee-Carter model to UK female mortality rates. Then the residuals from the model are plotted on the very top line against calendar year. The idea is that there should be no systematic pattern as you look across from year to year.

In the second one they are plotted against age. Again, there should not be a systematic pattern. It is a matter of judgment here, but they do not look too bad. Then they are plotted against year of birth and you can see we have waves. There is clearly a systematic pattern. That is picking up the fact that these data have a cohort effect. This is not a very good model for UK female mortality rates because the pattern there is a strong one. It shows the need for using the cohort effect, including that in the model. That is what we have done. We have also done it for males as well and for other data sets, assured lives, as I have mentioned.

The slide you can now see shows the effect. We have produced annuity values but in order to keep the number of slides under control we have only shown you projected mortality rates. These are for the year 2020. For the top three that you can see we have females mortality, UK; male mortality, UK; male assured lives. Similarly at the bottom.

In each of the panel is at the top of the slide there are three curves. The latest set of mortality rates, roundabout the year 2000; the age period, which is Lee-Carter forecasts; and then an age cohort, a model that just allows for age and cohort effects. At the bottom of the slide we have the latest, the Lee-Carter age period and then the age period cohort.

The striking thing is the way that these mortality rates fall at particular ages, the way that the cohort effect drags down mortality rates for 2020 at key ages. That is the cohort effect feeding through. So you get unusual patterns. You are not going to get the smooth, exponential shape that we are used to, and that is because of the cohort pattern coming through when one looks at future calendar years.

The slide you can now see shows the underlying parameters in the Lee-Carter cohort model, alpha, beta that multiplies the cohort effect, beta that multiplies the calendar year effect, kappa, the calendar year effect, showing the projections using the conventional time series methodology which does not pick up all the uncertainty.

Richard is going to elaborate on that. You can see the cohort effect with its projections.

[TAPE CHANGED AT THIS POINT]

... I think quite a lot. So when we read that years of birth 1920 to 1940 were particularly favoured in terms of future mortality, these cohorts were special. This is telling us that. A very sharp fall for these cohorts. It then flattens out. This is for females.

The corresponding male picture is even more dramatic. It shows you a heap, a mountain, and then almost flat afterwards.

If it is flat without many wobbles that is a good sign for projecting the future. But the hill is something which needs to be taken into account also for projecting the future. I think that this is a very important slide. It is a model that enables us to measure statistically the cohort effect.

In terms of development of the science, and this is really a side comment, we have two methodologies that are being worked on actively. One is the P-spine methodology that you heard about this morning. The second is the Lee-Carter methodology. A personal comment in a moment about the two. I have just received a report on my desk from Eilers, who is one of Ian Currie's co-workers, who really advocated the P-spline methodology for applying to mortality forecasting. In this latest report Eilers says that there is an advantage for actuarial applications in using Lee-Carter because it allows for flexibility. It fits the data and then does the forecasting in two separate operations. He thinks that there is an advantage.

But there are these two, so what should we do? My personal comment is that we should not abandon either. We need to keep using both because although I have indicated that there are some errors, hence the confidence intervals, and Richard is going to elaborate on those, underpinning all of this is model error. Without having more than one model, we are not going to get a good feel for model error. Even then the feel that we get for it is a little bit rudimentary.

I will hand over to Richard.

MR RICHARD VERRALL: I am going to talk a little bit about Bayesian models. There are is obviously a lot of uncertainty from a lot of different sources. We are concentrating this afternoon on mortality modelling but of course there is also the interest rate uncertainty as well. It is part of a large project that we have been looking at over the years to see how we can bring these together, include them all -- maybe it is a bit optimistic to do that -- and allow for all these sources of error when looking at the future.

I think that actuaries are getting much more used to looking at estimation error and process error, things like the ICAS regime have emphasised the importance of those types of things. In the past they would have concentrated much more on process error. Estimation error means that if you have a model you then have to estimate the parameters. You are not going to get the right answers. There is going to be estimation error and you have to take that into account. Even if you knew what the model was, you had the right parameters values and the right model, you would not be able to forecast the future because it is going to be a random outcome. There is process error there is well. Of course, you would not have the right model. They are only models anyway, so you cannot replicate exactly what is going to happen. Therefore you have to take into account the fact that you have model error.

I think that Bayesian models give you a very nice framework for including these. It may be a little bit of a different approach in that when we apply them we do everything all at once. We do not split it up into looking at the past data and estimating and then simulating for the future or something like that. It is all done at the same time. We have begun working on things whereby we can model the mortality of the past, project into the future, model investment returns and project those into the future and take it all into account at the same time.

I am going to concentrate on looking at mortality modelling. I will concentrate on looking at the basic Lee-Carter model. These are the advantages of Bayesian models: you can put in all the sources of error; it is quite common to include estimation error and process error. It is less common to incorporate model error as well, but I think people will work on that in the future. It is very flexible in two senses, one in that it allows you to look at whatever you want to look at. If you want to look at a particular portfolio or a group of lives or a whole population, a single life, a single risk or a particular policy, it is fairly straightforward to do that. It is a simulation based approach.

It also allows a lot of flexibility in the way that you do the modelling as well. I think that there are some advantages within the Bayesian modelling world which have not been emphasised so far.

As an illustration of how this might be used, I am going to talk about a Bayesian version of Lee-Carter. I am going to show you a few slides as to what you might get out of it. It is not the cohort model. It is just the basic Lee-Carter model with the log mortality rate being something which I think most people are familiar with now with the alphas, betas and kappas. I am not going to concentrate very much on the alphas and betas. I will just let those happen. I am more interested in kappa, where the time effect comes in.

The usual Lee-Carter approach is to use ARIMA type modelling. ARIMA is classical time series. You can fit it into a Bayesian model but it is more natural to do something slightly different, which is what I am showing on the slide on the screen.

It is made up of three parts. The first part says that kappa has a normal distribution with a certain mean, m t. I will concentrate on the means at the moment.

What is the mean? Well, the mean is the previous mean plus b t. It is basically a straight line model; b t is the gradient of the line. It just goes up in one time period by amount b t, usually called the growth in the Bayesian world. It is usually called a dynamic model.

A b t is also a random variable so the gradient does not stay constant. If we put all of those variances to 0 then we would get exactly a straight line and it would be a bit boring. What you can do is to choose different values for the sigmas to allow it to have more flexibility. One way of looking at this is to say locally it would be a straight line. Then it will follow the data bit and it can vary. Maybe when you look at it with a wider spacing of time it will not look so much like a straight line.

It is a fairly straightforward model. You can change the sigmas or try and estimate them from the data. Obviously, the larger you allow the sigmas to be, the greater will be the variation that you get there.

To give you an idea of the types of things that you can get out of this, we looked at a number of datasets. We concentrated on a couple first of all in order to make sure that when we fitted these models we got something sensible. We looked at what other people had done, something that Steve had done earlier with a colleague of ours, to make sure that the results looks sensible. Then we looked at what we were more interested in which was looking at pensioners' data.

I will show you a few slides and the talk about the types of things that you can get out of them.

Concentrating on cases A and B, those of you who have seen these types of pictures before, A is the grouped data, B is a different dataset, but not grouped. These are the alpha values and it is reassuring that from the Bayesian model where you are doing the estimation in a very different way you get very similar results for both alpha and in fact also for beta. These are similar to the results that Steve showed earlier on.

What is more interesting is what happens with kappa. We have this varying straight line. You get something which is very similar to the patterns that you get out of a straightforward Lee-Carter model. In these cases you can see the change in the gradient happening later on. It has been allowed to happen quite naturally because the gradient has been allowed to vary through the stochastic model.

It all works very nicely and it can be used to model Lee-Carter. When you come to do the forecasting, it will project forward a straight line because that is the underlying model. Obviously, there are decisions to be made there. Given what has happened at the past, given what the data are telling you, this is what the Bayesian version of Lee-Carter will tell you to do for the future.

As I said, what you can do is to build on top of that whatever you want. Here, what I have done is to look at a single female life age 60, an immediate annuity, with a 6% interest rate and to look at the distribution of the value of the annuity for that single life. If the life dies early, the cost will be less. It drops off a little too quickly for my liking, which is because I have not allowed lives to live beyond a certain age. Maybe I could allow that a little bit more than at present. It cuts off at a single age.

What is happening here is all at the same time the data is being used to estimate the parameters. The mortality rates for the future are being simulated using those parameters. Therefore you get the estimation error coming in, then the amount of time that a life lives for is being simulated on top of that. Therefore you get the process error coming in. What we have also done is to allow a stochastic model to be built for the interest rate as well to allow that uncertainty to be projected into the future. What we are also looking at is including a number of different scenarios or models so that you can include some sort of model error in this as well.

What I have shown you here is something quite simple. It is a single life. It is possible to do portfolios or sets of lives, different things, different types of lives, and so on.

Questions

THE CHAIRMAN: Steven, could I ask you a question? You mentioned the extra flexibility that comes from fitting separately from projecting, from forecasting. Could you expand on that in terms of the benefits of that extra flexibility? It was not immediately obvious to me how one would use that.

PROF HABERMAN: I was quoting somebody else. I think what they meant is the following. In the Lee-Carter structure, one fits a set of parameters and there is quite a bit of flexibility within that in terms of, for example, one can smooth the alphas; one can smooth the betas. One distils from that the kappas, which are the key thing for understanding the time trend. Then one separately forecasts the kappas.

One could do the same if one fitted Gompertz's model to a series of calendar years data, obtained the key parameters and then tried to model the trend in those parameters. What Lee-Carter seems to have about it is that the forecasts you get from doing that method are more satisfactory in terms of their shape and feel than the ones that you get from fitting cross-sectional data using, for example, Gompertz, Makeham or some Heligmann[?] and Pollard[?] complicated formula.

I think that that is where the flexibility comes in, because it is two-stage, because there are some extra components of flexibility that you can build into the first stage in terms of the alphas and the betas and because of this focus on extracting the key driver, the key time effect driver.

Mr [?] (Hannover Life Re): I find these sorts of exercises can sometimes be dangerous. Often I have seen attempts, coming from a pricing background, to quantify future mortality improvements. This is more of a comment than anything else. I am frequently presented with a tabular form of A/E. It will say "This is what the level of mortality was in 1984. This is what it was in 1985, 1986, and so on, up to today. So we obviously want to guess, as to what the future levels of mortality are going to be using this because we have no better ideas, so were going to use P-splines." Then the moment they start drawing these graphs and having random variables, driving random variables, driving the parameters of random variables, and crazy things happening, people's eyes glaze over. When it is all finished they go "Okay, so it is 1.5%. Right, we will go with that."

In the whole process no one pays attention to what constitutes the basic data that is driving these things off. That creates a false sense of security. I guess that we all know this. It often feels like overkill.

With the Bayesian model, for example, as many merits as it may have, why stop there? Why not say "Okay to account for model error, we should have a Markov Chain and we should switch between the Bayesian model and the Lee-Carter one to allow a for model error. Let us have another model so we will have a random model that switches between." It is a straight line!

I often feel a person gets sucked into the detail of it and because people feel that they cannot understand it they suddenly have more confidence than they should have.

THE CHAIRMAN: That was a comment rather than a question, so you do not have to respond.

PROF HABERMAN: I feel I should like to respond. I do not disagree with you, but I do not think that this is modelling for the sake of it. This is not building a structure just because it is nice to build a structure and add bells and whistles. We are all interested in this because it is an important problem, and it has not been cracked to produce reliable forecasts.

Investigating new models and adding extra features because they are needed is really the motivation here. I agree, we need to keep our feet on the ground. We also need to remember the reliability of the data that we are using.

Mr [?] (Hannover Life Re): My key gripe is: present this to somebody who has not seen a random variable for five years, they have fundamentally lost touch with what statistics is all about. They will suddenly look at this and feel intellectually incapable and they will say, "I cannot remember why we are actually looking at this, but I am going to believe its results.” That is something that worries me. They frequently do not come with sufficient health warnings.

If somebody comes to me and says, "I want to know what you think about future mortality" I give them an interactive answer and I say “Here is the spreadsheet. This yellow block is the key assumption. The other things are in the background which may or may not make sense.” Fiddle with this thing and it will really give you an idea how little we know in this situation.

PROF ANGUS MACDONALD: I have an observation on the last comment. If you gave a modern FSA return to someone who had not looked at technical life office matters for the last ten years and they said "I do not remember this, but I would just go with it", you would take away their practising certificate.

MR MATTHEW EDWARDS: This is a question in particular relation to the current graph shown by Professor Verrall. I was particularly interested in the last presentation getting more of a grip on the variability, the uncertainty, surrounding all these things. My question is with the current graph I was not entirely clear whether you were parameterising everything and then doing lots of stochastic runs using those parameters, or were you doing many different parameterisations of the data?

Then a related question which may have a bearing on that is: in your previous formula where you had a succession of formulae each with their own variants, presumably the extent to which you decide on that variance parameter has a huge impact on the overall variability of your results. It was not clear to me whether you were fitting those variance parameters or tuning them *a priori*.

PROFESSOR VERRALL: Looking at the graph, I suppose that the easiest way to explain it is to say that we use parameter values taken from the distribution of parameters. Using those, we then simulate the future lifetime. So it includes the estimation error on the parameters.

On whether the choice of variances can affect it, for this run we did it with what we thought were sensible choices for those parameters. We have not estimated from the data yet. We will be doing that. Relating to what you just said, the input of so-called expert opinion into this type of modelling is something which could be looked at as well, not only just to change a parameter but to try to find a sensible way of encapsulating the uncertainty that you have about what that parameter should be, and including that within the model is something which could be done.

It is not meant to be magic. It is meant to be helpful in order to bring these things in and to do something which you think is sensible.

[?]: I have a quick question. This is all cracking academic material. I will not hold you to the answer to this question outside of this room, but just for those of us who have to face clients, what is your best estimate of the life expectancy of a male currently aged 60 and what will the same answer be in 20 years' time?

PROF HABERMAN: May I reply in writing?

[?]: I would not if I were you!

PROF HABERMAN: But I do not think it is academic!

Chairman's Closing Remarks

THE CHAIRMAN: If I could just close the seminar, frequently we attend seminars in the hope of getting answers. I suspect today that we may have actually raised more questions in people's minds than we have answered. But at least, thanks to the presentation of John Lister and Nigel Bodie, they have helped us to know how to communicate that increased level of uncertainty when we get back to the office.

The discussions that we have had around mortality and the projection methodologies were certainly not intended to be any form of big brother exercise whereby methodologies were voted out of the house. Steven summarised it in terms of we do need to keep using various methodologies to see what insights they give us because obviously it will only be in 50 years' time that we can actually decide which one was the best for that period. It certainly may not be the case that it is the best going forward from there.

The key question for the profession is probably how we take this forward in a collective way and in a sensible manner.

Something I absolutely hate at the end of training courses and the like is when you are asked to come up with three action points or three key actions. But I have written down three for me today.

Thanks to Brian Wilson, I know that I need to invest more in my pension plan in order to improve my life expectancy! Richard Willets suggested that perhaps I do not need to be quite so concerned as I was about my own personal increasing obesity. Nigel Bodie told me to move to Penzance -- in a very polite way.

Could we finish by thanking all of the speakers today and I should also like to thank Hannah Bolton who joined the actuarial profession just in time to get lumbered with organising this seminar and has done a wonderful job. So

I should like you to express your thanks to all of the speakers and to Hannah. [Applause]
