Continuous Mortality Investigation

Mortality Committee

Revised Working Paper 20

Stochastic projection methodologies: Further progress and P-Spline model features, example results and implications

IMPORTANT NOTE:
Working Paper 20 was originally published in April 2006. This revised version was published in November 2007 (together with a revised version of Working Paper 25).

A description of the issue relating to the exposure data that gave rise to the amendments in this Working Paper is contained in a separate document entitled "Errata to CMI Working Papers 20, 25 and 27 on Mortality Projections".

This version of Working Paper 20 has NOT been updated for any developments subsequent to its original publication, other than the noted errata.

November 2007

Continuous Mortality Investigation

Mortality Committee

Working Paper 20

Stochastic projection methodologies: Further progress and P-Spline model features, example results and implications

1. Introduction

1.1. Since the ‘a(55)’ tables were published in 1953, the CMI has customarily provided mortality projections – in some form – when it has published mortality tables. As recently as the “92” Series tables, the CMI produced a single set of projections, derived by considering past trends within both the CMI’s own experience and population mortality data.

1.2. The most recent set of projections were presented in Working Paper 1 (CMI, 2002) which reviewed historic methods, and then briefly introduced the new projections. These were distinctive in three important respects.

a) They recognised the so-called cohort effect; that is, the dependence of mortality improvement rates on a person's year of birth. See Willets (1999) and Willets (2004) for more details.

b) These projections were ad-hoc extrapolations based on the results of a methodology new to actuaries, namely penalised spline regression (or P-Splines). See Working Paper 3 (CMI, 2004) and the references therein for more details.

c) Three alternative projections were offered, instead of the traditional single projection. This was an explicit recognition of the uncertainty of long-term mortality projections, although it was not a probabilistic statement.

1.3. The three scenarios published in Working Paper 1 were called ‘short’, ‘medium’ and ‘long’ cohorts, differing in the length of time over which the cohort effect was assumed to persist. These periods were chosen arbitrarily and no probabilistic interpretation was possible. They were described as ‘interim’ projections, signalling to the profession the CMI’s intention to undertake further work in this field. As well as being consistent with the cohort effect, the use of probabilistic models for mortality projections would be consistent with the trend in the regulatory environment of placing greater emphasis on risk management based on stochastic models of risk.

1.4. As a result the Mortality Projections Working Party was established to explore possible methodologies and develop these for use with the “00” Series tables. Much work has been undertaken towards that goal and this paper seeks to provide an overview of that work in Section 2, up to and including the publication of Working Paper 15 (CMI, 2005c) which proposed two possible models for use: P-Spline and Lee-Carter. Section 3 then provides a summary of the feedback received to date on Working Paper 15, together with the Working Party's responses. Section 4 reviews the CMI's most recent activity, in particular relating to illustrative software developed, and workshops held in mid-September in London and Edinburgh and the presentation of initial results at CILA in late September.

1.5. Simultaneously with the publication of this paper, a revised version of the software is being made available. The P-Spline functionality has been amended in the light of comments received on the original version and we hope will be easier to use, however users should be aware that the software is still only intended to be illustrative. There have been some delays in finalising the software for the Lee-Carter model. Due to the computer-intensive nature of this model, it is taking a long time to produce and analyse results and the Working Party is still considering the results. Therefore, the features of the Lee-Carter methodology and sample results will be covered in a separate Working Paper which the Working Party expects to publish in the second quarter of 2006. The Lee-Carter functionality is disabled in the March 2006 version of the software.

1.6. Subsequent sections of this paper provide the background to examples based on the P-Spline methodology. Section 5 briefly describes the models and datasets used. Section 6 then describes how to use the output from the P-Spline modelling software made available by the CMI. Section 7 discusses various features of the P-Spline models. Section 8 provides sample annuity values for use in 2004 produced using the P-Spline projections and data to 2003 which are then compared to annuity values based on the interim cohort projections. This section also discusses how the equivalent sample annuity values for use in each year between 1993 and 2004 would have progressed as more data became available over the period.

1.7. Section 9 sets out tentative conclusions regarding the P-Spline methodology and sets out the next steps. Clearly, much further work is required but that is beyond the scope of the Working Party. This further work, described in Working Paper 3, would include consideration of questions on model uncertainty, the correlation between mortality and investment risk and moving the projection methodology towards cause-specific projections. However, better data on cause-specific mortality may be required before there can be much further progress on cause-specific projections.

1.8. This working paper has been prepared for the Mortality Committee of the CMI by a Working Party consisting of Angus Macdonald, Adrian Gallop, Keith Miller, Stephen Richards, Rajeev Shah and Richard Willets. It has been approved by the Committee.

2. Summary of previous work

Background

2.1. The exploration and development of possible mortality projection methodologies was a major piece of new work, and the first step the Working Party took was to host a joint seminar on mortality projection with the GAD in October 2003. The purpose of this seminar was to hear speakers from outside the actuarial profession. The second step was to produce a discussion paper on the issues, which was Working Paper 3. The third step was to consult the profession on these issues, which was done by asking questions at the end of Working Paper 3 followed by a meeting in Staple Inn Hall in June 2004. The responses made both at the meeting and in subsequent private letters were summarised in Working Paper 11 (CMI, 2005a). The fourth step was that the CMI sponsored a CASE award, in conjunction with the Engineering and Physical Sciences Research Council, to support a PhD project at Heriot-Watt University, under the supervision of Dr Iain Currie, into smoothing and projection methodologies for longevity projection. The fifth step was that the CMI supported a two-day technical workshop on mortality projection in Edinburgh in September 2004, organised by Dr Currie. The next step in this process (in Working Paper 15) was to describe work to date and to indicate modelling frameworks for stochastic mortality projection that the Working Party feels warrant further consideration. Finally software has been developed by James Kirkby and Iain Currie to illustrate the methodologies. Workshops were held in London and Edinburgh in mid-September 2005 to discuss the methodologies and demonstrate the software.

2.2. These steps are described in more detail in the remainder of this section; however for full details please see the original papers and the references therein.

Summary of Working Paper 3

2.3. Working Paper 3 was a discussion paper. It described the basis and background to the cohort effect, and the interim cohort-based projections which resulted in Working Paper 1. Working Paper 3 also contained a description of a joint seminar (with the GAD) held in Edinburgh on 6 October 2003 to discuss the views and approaches of demographers, statisticians and gerontologists, all of whom have a strong professional interest in the projection of future mortality and its underlying causes.

2.4. Working Paper 3 described the background to the need for a new set of projections. This included the need to give some indication of uncertainty in projections, and to allow a more transparent approach to risk management. These are required as insurers are increasingly moving, or being moved, towards the use of such risk management tools, not least because of the IASB `fair valuation' rules, FSA `realistic balance sheet' requirements, and convergence of regulatory regimes in banking and insurance.

2.5. An overview of different projection methodologies was given in Section 2 of Working Paper 3, contrasting their strengths and weaknesses. Process-based projections attempt to model trends in causes of death, although this approach is not favoured because of problems in death classification and insufficient understanding of the major cause-of-death processes. Extrapolative methods are based on projecting historical trends in mortality into the future, although all such methods include some element of subjective judgement, for example in the choice of period over which the trends are to be determined. Examples of different approaches to extrapolation are discussed in Section 2.1 of Working Paper 3.

2.6. Section 2.2 of Working Paper 3 considered the various types of model in use, including the current CMI methodology, and the methodology used by the GAD for projecting mortality in the official national population projections for the U.K. and its constituent countries.

2.7. Special consideration was given to the treatment of uncertainty about future projections. In particular in Section 4, three well-known sources of uncertainty associated with the use of statistical models were discussed:

a) model uncertainty;
b) parameter uncertainty; and
c) stochastic uncertainty.

2.8. Working Paper 3 also discussed the issues associated with fitting data and making projections. The first major question is what data set to use, and there are particular problems associated with the lack of suitable annuitant mortality experiences. Considerable care is required with projections, particularly where the experience of an annuitant population of insufficient size could potentially lead to implausibly narrow confidence intervals in projections. This problem was illustrated by an example where there were insufficient data to refute the fitting of a straight line to the mortality trend, which, by virtue of the model structure, resulted in very narrow confidence intervals around the projection. Another issue specific to regression models was that traditional polynomial methods can yield acceptable fits in the region of the data, and yet produce very poor projections outside it.

2.9. Working Paper 3 contains a lot of detailed, technical discussion and is a useful backdrop to understanding subsequent work. It also posed explicit questions for the profession and details of the responses were published in Working Paper 11.

Summary of Working Paper 11

2.10. Responses to the questions posed in Working Paper 3 were received during a seminar at Staple Inn on 4 June 2004, and also in writing from some life offices and reinsurers. Full details of the questions and responses are contained in Working Paper 11. The elements relevant to this paper are set out below.

2.11. The second question concerned the appropriate level of aggregation in projecting future mortality. There was general agreement that, while cause-specific projections should be carried out if they could be made to work, the currently available methods and data fell far short of being adequate. However, contributors at the seminar expressed their continuing concerns that the improvements seen in the past may have been greatly influenced by changing smoking patterns and, as smoking patterns stabilised, future improvements could follow a different pattern. The Working Party agrees that changes in smoking incidence is indeed an important factor affecting mortality improvements, and one of those factors most easily identifiable from the available data. However, the Working Party does not believe that changes in smoking incidence explain all patterns; see Willets (2004). Mortality improvements for almost all causes arise due to the interaction of a number of factors. Many of these factors are common to several causes and few of the interactions are fully understood. Therefore, it is very difficult to model cause-specific mortality rates in a robust way.

2.12. The third question was whether the CMI should continue to project cohorts. All respondents supported the projection of cohort mortality improvements, giving as their reasons the evidence of the presence of cohorts in past mortality improvements.

2.13. The fourth question was whether quantitative measures of uncertainty associated with projections were needed, and, if so, what form they should take. Feedback was unanimous on the need for a measure of uncertainty, although views differed on how this should be provided. Scenarios similar to the long, medium and short interim cohort projections were seen as very useful in presenting the mortality risks to non-actuaries, particularly boards of life offices. Some respondents indicated that they would like stress-testing scenarios to be provided. In an informal show of hands, attendees at the seminar voted overwhelmingly in favour of measures of uncertainty being provided with the next set of projections. No one voted against.

2.14. The fifth question was whether distributions or percentiles of future rates of mortality, derived from statistical models of past rates of mortality, were sufficiently meaningful to be used in practice. There was general agreement on the need for quantitative measures.

2.15. The sixth question was whether projections and any measures of uncertainty should be based on the largest available appropriate populations. The feedback agreed that the largest appropriate population should be used and that this choice was greatly affected by what was practical. The Working Party believes that the only choice is between population data and male assured lives data. No other current CMI experience is old enough or large enough to be credibly used for producing projections.

2.16. The seventh question was whether there was any clearly preferred methodology. The majority of the respondents had no preference for any particular methodology, although a desire for simplicity was expressed. Concerns were aired about the possible complexity of stochastic mortality models and the difficulties that would arise in explaining them to non-actuaries. The Working Party did not expect the mortality model to be too sophisticated or complicated, compared to stochastic asset models for instance, and so such difficulties should not be overplayed.

2.17. The eighth and final question was what the financial consequences of allowing for uncertainty in projecting future mortality might be. There was concern about how mortality research could be misunderstood, especially outside the profession. In particular attention was drawn to the possibility of an investment analyst concluding that offices were going to strengthen reserves, with a consequent impact on statements of earnings or profits. There were also concerns that regulators may draw inappropriate conclusions from the results of research. The need for careful communication and appropriate caveats was highlighted.

2.18. In its subsequent discussions, the Working Party rapidly concluded that it could not possibly produce any definitive answers to the problems raised by mortality projections, and that any methodology that it might suggest for use in practice would inevitably be subject to criticism and to change in the light of ongoing research. We believe it is absolutely essential that users of the projections are fully aware of this.

2.19. In particular, the Working Party considers the following questions to be beyond the scope of our current research:

a) model uncertainty;
b) correlation between mortality and investment risk; and
c) moving projection methodology towards cause-specific projections.

2.20. The Working Party also noted that any financial uncertainty arising from uncertainty regarding the level of aggregate future mortality rates can be swamped by the heterogeneity of the amounts of pensions within an office's portfolio, given the very large difference between the mortality of pensioners with the smallest and largest pensions. For smaller portfolios, (indicated by our modelling in Section 3 of Working Paper 15 to be fewer than 5,000 lives), the heterogeneity in pension size can be one of the biggest drivers of financial uncertainty, especially over shorter time-periods.

Summary of Working Paper 15

2.21. Working Paper 15, published in July 2005, contained an update on the CMI Mortality Projections Working Party’s work towards developing stochastic methodologies for projecting mortality in association with the “00” Series tables.

2.22. Working Paper 15 drew comparisons with the use of stochastic methods in Asset-Liability modelling, with which actuaries are increasingly familiar. It then set out the P-Spline and Lee-Carter models in detail. Section 2.4 compared the two approaches whilst making clear that in the Working Party's view, neither could be regarded as superior.

2.23. Section 3 provided examples of the risk-capital requirements for model portfolios of varying sizes derived by fitting a P-Spline model with penalties on age and period to male population data. These examples encompassed both the capital required to cover the systemic risk, arising from the underlying data, and the stochastic uncertainty facing smaller portfolios. These examples were based on population data to 1990 and were intended to illustrate principles rather than provide any form of guidance.

2.24. Working Paper 15 concluded with five specific questions to which responses were requested:

1. Are potential users in favour of the broad approach of developing a stochastic methodology?
2. Is there a preference for Lee-Carter or P-Spline models, or should both models be made available?
3. What are users' views on the possible production of spreadsheet models utilising Excel and R, and/or the possible production of a CD containing a suitable number of scenarios, and/or any other form of output that may be desired?
4. Should the models allow users to make appropriate adjustments to the projections (say to reflect the socio-economic class mix of their business) or should the implementation of such adjustments be left to the discretion of individual users?
5. Should the Working Party specify a preferred basis and methodology, or should this be left to the discretion of individual actuaries?

The responses received are considered in the following section.

3. Feedback on WP 15

Summary of responses

3.1. Responses to Working Paper 15 were requested by 31 August. Nine responses were received – from Deloitte, GE Insurance Solutions, Hibernian (Ireland), Legal & General, Norwich Union, Pearl, PwC, Scottish Widows and Standard Life. The CMI would like to thank those who took the time to respond.

3.2. At a very high-level, the responses to the specific questions can be summarised as follows:

- There is broad support for the introduction of stochastic methodologies;

- Most respondents felt both P-Spline and Lee-Carter should be made available - where a preference was expressed, there was no clear ‘winner’;
- There was broad support for the CMI making illustrative software available, but some concern over the use of R;
- It was clear that respondents felt that the software should allow actuaries to make appropriate adjustments;
- Respondents clearly felt that it was inappropriate for the CMI to prescribe a method or basis;
- However there was an equally clear demand for some guidance, both to assist actuaries in gaining understanding of a new area and to avoid inappropriate discrepancies between companies’ approaches.

Responses to the specific questions

3.3. Of the nine responses, seven responded to the specific questions (whilst the other two made generic responses only).

Q1. Are potential users in favour of the broad approach of developing a stochastic methodology?

3.4. The majority of responses answered “Yes” although there were a few qualifications to this, namely that the projections and associated measures of uncertainty are meaningful and their status in accordance with professional guidance is clear. The point was also made that liability valuation systems generally require deterministic mortality assumptions so the projections must enable deterministic assumptions to be chosen from the available stochastic output.

3.5. There was one exception, where an office is not convinced that a stochastic approach is necessarily appropriate for modelling future mortality trends. They have already developed a different approach for their ICA calculations based on projecting the impact of significant health and medical improvements on mortality rates.

3.6. The CMI is not presenting stochastic methodology as the only solution in this area and welcomes contributions to the debate. We are very keen to learn more of this approach and to compare the outcomes with those of the stochastic methodologies under evaluation.

Q2. Is there a preference for Lee-Carter or P-Spline models, or should both models be made available?

3.7. All of those responding specifically to this question agreed that both models should be made available. Several responses acknowledged that both appear to have merits, whilst others expressed a preference for either Lee-Carter or P-Spline.

3.8. One response commented that although WP15 criticises the Lee-Carter model for not explicitly allowing for cohort effects, this was now possible and referred to Renshaw & Haberman (2005). This is an area of current work that had not previously been brought to the attention of the Working Party. We are grateful to this respondent and to Professor Haberman for subsequently sending an updated version of the paper which we understand is expected to be published in 2006. The Working Party has not yet had the opportunity to consider this work.

3.9. Another response commented that there is a need to project future mortality tables, as done by the Olivier-Smith model (see Olivier & Jeffery (2004)). This is not an area the CMI has yet considered in any detail, but suggests that further research is required.

Q3. What are users' views on the possible production of spreadsheet models utilising Excel and R, and/or the possible production of a CD containing a suitable number of scenarios, and/or any other form of output that may be desired?

3.10. This was supported – to a degree – by all respondents. Concerns expressed included:

- The use of a package (R) that is not commonly used within the insurance industry, because of opposition from IT and lack of transparency;
- Concern that users may confuse the selection of best estimate mortality with the estimation of variance around the best estimate;
- Concern was expressed that the provision of parameter runs on a CD might mean that companies don't fully understand the models and the implications of varying the parameters.

Whilst the CMI understands the concerns expressed about R, it is widely used in academic and other circles and increasingly within the insurance community. It has been used for the initial software developed by the CMI (see Section 4 for more details). The other concerns both appear to relate to possible mis-use of the CMI software. It is certainly not the CMI's intention to provide a "black box" that actuaries or others should use without understanding the context. We hope that users will have taken the time to read Working Paper 15 and other relevant papers to understand the issues that we hope the software sheds some light on.

Q4. Should the models allow users to make appropriate adjustments to the projections (say to reflect the socio-economic class mix of their business) or should the implementation of such adjustments be left to the discretion of individual users?

3.11. There was general support for this flexibility, although one response appeared to suggest that they were looking to the CMI to provide the different assumption sets, which was not the Working Party's intention.

Q5. Should the Working Party specify a preferred basis and methodology, or should this be left to the discretion of individual actuaries?

3.12. The answer to this question was clearly that the CMI should not be prescriptive, however there was also a clear demand for some guidance.

3.13. One response suggested that the CMI's work should facilitate a consistent approach with differences being valid (i.e. reflecting only past experience, market segmentation etc in some consistent way). Other responses echoed this with references to "… provide a common starting point…" and "… to avoid arbitrary differences across the industry…" One response suggested that the CMI provide a 'prudent' set of factors which individual users could reduce subject to detailed analysis on their part, drawing an analogy with the Guidance Note on market risk for Individual Capital Assessments. The issues facing offices that only have small blocks of business, and hence no statistical credibility to their own experience, were also noted.

3.14. This is clearly a difficult area, and was recognised as such by several respondents. Some of the views seemed to contradict themselves, for example "We would like the Working Party to provide guidance on a preferred basis and methodology without being specific on a single preferred basis." Several comments referred to the role of the profession with regard to setting standards.

3.15. Overall, the responses received appeared to suggest that the CMI should go further than had been planned in providing guidance. Whilst recognising this feedback, the Working Party is strongly of the opinion that the methodologies require further evaluation within the profession as a whole before all the issues can be identified and any clear guidance issued by the profession.

Additional points

3.16. In addition to the responses to the specific questions, respondents made a number of other points, including:

a) One respondent commented that Working Paper 15 ignored issues for projecting trends in assurance mortality. In particular there was a concern that the models being evaluated may smooth out too much of the "roughness" that is important in this context.

This was akin to a comment that the most appropriate model may depend on the purpose – e.g. pricing, reserving and setting capital requirements might have different needs.

The Working Party accepts that Working Paper 15 was focused on annuity business, and that the need to model fluctuations in experience needs to be considered in further work. In this regard Lee-Carter may be more appropriate for work on mortality risks than P-Spline. We also certainly accept the need for the method used to be appropriate to the purpose and feel that it is only through wider evaluation of the methodologies that this the relative benefits will be fully recognised.

b) One respondent asked whether improvements and variability are connected. Reference was made to Olivier & Jeffery (2004), which suggested that England & Wales population mortality showed two distinct phases, the later of which had faster improvement associated with reduced year by year variability.

This is an area that the Working Party hopes will be researched further.

c) One respondent stated that they are keen that the "00" Series base tables and projections are published as soon as possible and therefore would not welcome at this stage any changes of direction or further investigations that would create further delay. However, they support the undertaking in the future of the further research mentioned in Section 5.6 of Working Paper 15.

Another respondent stressed that enough time should be spent researching and understanding the methods used and the results that arise before an approach is prescribed by the FSA or the actuarial profession. They also commented on the need to educate actuaries and non-actuaries (such as company boards) on the new concepts.

The CMI hopes that the first comment is addressed in the progress being made on the base tables whilst research into projection methodologies continues. We also agree with the second comment. It is important to note that although the CMI is evaluating these two methodologies for mortality projections, it will not be seeking approval for either methodology nor does it rule out alternative approaches. It is only by exposing the work fully to the profession that the benefits and features of the methodologies will become apparent.

d) One respondent commented that other models need to be considered, specifically referring to the Olivier-Smith model and Yang (2001).

Working Paper 15 stressed that other models need to be considered and the CMI intends to continue its research but hopes that this will be broadly supported within the profession.

e) One respondent commented on the importance of projecting mortality tables rather than mortality rates, as it is the following year's mortality basis that will drive capital requirements. They felt the Olivier-Smith model was a step in this direction.

The Working Party does not believe that the Olivier-Smith model has yet been subjected to rigorous external appraisal, but feels it certainly warrants further review, particularly for the circumstances for which it was developed. The ultimate goal would appear to be a need for a "Game Theory"-type approach to predicting how actuaries will react to developments in setting future bases. This is outside the scope of the CMI's current work.

f) One response highlighted what were felt to be significant weaknesses with Lee-Carter models. In particular that they do not allow the user to fit an initial mortality expectation curve, perhaps leading to confusion between the initial base table and subsequent random variations, in the context of ICAs.

The Working Party has not yet had an opportunity to review or discuss these comments in detail but hopes to do so before publishing the Working Paper on the use of Lee-Carter methodologies referred to in Section 1.5.

4. Subsequent activity

4.1. Since Working Paper 15 was published in July 2005, the CMI has released illustrative software, hosted workshops in London and Edinburgh and presented results at the Current Issues in Life Assurance (CILA) seminar. These are briefly described in this section.

Software

4.2. The software was made available with the aims of illustrating the P-Spline and Lee-Carter methodologies, assisting the profession to debate the issues and to help users validate the results of their own work. The CMI would like to express its thanks to James Kirkby and Dr Iain Currie of Heriot-Watt University for developing this software.

4.3. The software allows users to fit and project using P-Spline and Lee-Carter methodologies. It is available via the CMI's pages on the profession's website. Simultaneously with the publication of this paper, a revised version of the software is being made available (see Section 1.5 for more details).

Workshops

4.4. These were held in London on 12 September 2005 and in Edinburgh on 14 September. During the morning both the P-Spline and Lee-Carter methodologies were discussed and the software was demonstrated for each. Illustrative results were then presented using output from both methodologies. These results used only data up to 1992. The results were used to illustrate:

- Model differences: P-Spline (period penalty) v P-Spline (cohort penalty) v Lee–Carter v "92" Series;
- the effect of using different datasets for P-Spline;
- the effect of using different parameters for P-Spline;
- the progression of projections for P-Spline as each additional year's data becomes available from 1984 to 1992; and
- the similar progression of annuity values for P-Spline from 1984 to 1992.

4.5. The morning concluded with a short discussion session on “Choosing a mortality basis for reporting purposes”, essentially considering some of the practical issues involved in using stochastic methodologies in a life office or pension scheme.

4.6. The afternoon provided an opportunity for attendees to experiment with the software and discuss issues with members of the Mortality Projections Working Party

4.7. Copies of the slides are available via the CMI’s pages on the profession’s website.

Current Issues in Life Assurance (CILA) Seminar

4.8. The CMI presented an update on the proposed “00” tables and the stochastic longevity projections at CILA in London on 28 September 2005. The results presented used output from P-Spline methodologies and illustrated the effect of using different datasets and parameters, similarly to the earlier Workshops. However, this time, results were based on CMI and ONS data to 2003 and the progression of projections and annuity values were shown from 1984. The results presented at CILA did not include any results using output from the Lee-Carter methodology.

4.9. The CMI had investigated the use of P-Spline methodology to model some non-UK mortality experiences and these results were also presented in the form of contour maps. The contour maps showed that cohort effects could be seen in the mortality experiences of some other countries in Western Europe and were not unique to the UK. Copies of the slides are again available via the CMI’s pages on the profession’s website.

5. Models and datasets considered

5.1. In this paper we show the results using the P-Spline models based on the latest available datasets, which are the CMI assured lives data from 1947 to 2003 and the ONS England & Wales population data from 1961 to 2003. The models are fitted to these datasets and then used to project scenarios of future mortality improvements. The resulting projected improvement factors are then combined with base tables of q_x to produce annuity values. The base tables used are the “92” Series[1] and also the proposed “00” Series[2]. The ONS datasets used by the CMI are not yet publicly available and were only made available to the CMI for research purposes. Therefore, the CMI can only show the results from modelling these datasets and cannot provide these datasets to others.

[1] CMIR 16 and CMIR 17
[2] Working paper 16 from the CMI

5.2. As described in Working Papers 11 and 15, the P-Spline model can be fitted to mortality data using either age-period or age-cohort penalties[3]. We show example results using both approaches and comment on their features. Whilst the P-Spline model can also be fitted using cohort-period penalties, it is not covered in this paper as, for the purpose of calculating future mortality rates, the Working Party feels that it produces poor results since it ignores the age structure that dominates in the mortality data.

5.3. For our purposes, the datasets needed to model future mortality improvements must contain age specific data for successive years and cover a much longer period than is necessary to simply graduate mortality rates. We consider that a minimum such period is 20 years. Additionally, for the age-ranges fitted, a large amount of data is required in each year of observation. The only UK datasets available to the Working Party which fit these criteria are the ONS England & Wales population and the CMI Assured Lives datasets which all cover "lives" only.

5.4. The ONS England and Wales datasets relate to the mortality experience from 1961 to 2003 and cover ages 0 to 100. In the ONS data made available to the CMI, for both males and females, deaths in each calendar year are grouped by age last birthday. The exposures are mid-year estimates based on decennial census information but adjusted for each subsequent calendar year to allow for births, deaths and estimated net migration. However, as data for ages above 89 are aggregated in this dataset, we have ignored data at these ages. We have also ignored data below age 20 because the steep falls in the mortality rates at the very young ages skew the P-Spline fits.

5.5. For males, the CMI Assured Lives dataset relates to experience from 1947 to 2003 and covers ages 20 to 100. This dataset is based on exposures and deaths, grouped by the nearest age. Deaths are based on the calendar year of death and include those notified to life offices up to 6 months after the year end. Exposures for each calendar year are based on in force policies at the start and end of the year, adjusted by the number of deaths. In this dataset observed crude death rates often decrease at ages above 90. Thus data at these ages are not, for the purpose of this paper, considered reliable and have not been used in the main results reported herein.

5.6. For females, the CMI Assured Lives dataset relates to experience from 1983 to 2003. Although the CMI started collecting data for this investigation from 1975, data prior to 1983 is only available in the form of aggregated age bands and quadrennia. With annual exposures of less than 10,000 lives at ages above 65 and less than 5,000 lives above age 70, data volumes are much lower in this dataset compared to that for males. Given these small data volumes,

[3] See Richards, Kirkby and Currie (2005) for an explanation of how the complexity of P-Splines actually draws on a relatively straightforward base. This paper also demonstrates practical issues arising from the use of age-period or age-cohort penalties.

particularly at the ages crucial for projecting longevity of pensioners, we have not provided any example results based on this dataset and do not comment on any other features of the dataset.

6. How to use the P-Spline model output

6.1. The P-Spline models fit forces of mortality (i.e. μ_x) to the data. As all the CMI datasets are adjusted to provide exposures and deaths measured at the nearest age, the fitted sets of μ_x based on this data apply to ages x exact. The models produce values for $\mu_{x,t}$ at each age x and for each year t within the fitted region of the dataset and in the region of the projection.

6.2. The updated P-Spline software produces values for the log mean values of $\mu_{x,t}$ at each age x and for each year t within the fitted region of the dataset and in the region of the projection as well as estimates of the standard deviations of the log mean values of $\mu_{x,t}$, which are denoted by $\hat{s}_{x,t}$. Scenarios can then be generated by sampling from a standard normal variable Z and applying them to the standard deviations to adjust the mean values of $\mu_{x,t}$ in scenario i. Thus each set of adjusted $\mu^{(i)}$ is produced from

$$\mu^{(i)}{}_{x,t} = \exp\{\log(\mu_{x,t}) + Z \times \hat{s}_{x,t} \}$$

Thus the set of adjusted $\mu^{(i)}$ relating to a particular percentile can be calculated directly by using the Z value for the percentile (e.g. using a Z value of 1.96 for the 97.5 percentile).

6.3. Fitting the models to the ONS datasets provides values for $\mu_{x+½,t}$ and $\hat{s}_{x+½,t}$.

6.4. As explained in Working Paper 15, the P-Spline model generates percentiles rather than sample paths. Where the P-Spline output is being compared to the results from other models that produce sample paths, care is required to ensure that percentiles and sample paths are considered consistently. For example, if 100 sample paths were generated, the 95th percentile would be constructed by choosing, for each projected year, the mortality rate corresponding to the 95th highest value, across all 100 sample paths, in that year. That is, 95% of the projected mortality rates are lower than those given by the 95th percentile. Note that 95th percentile annuity values (i.e. 95% of the projected annuity values are lower than these amounts) are generated from 5th percentile mortality rates.

6.5. Note also that basing annuity values on mortality rates from the 5th percentile gives higher values than those calculated from the mortality rates of each of the 100 sample paths and selecting the 95th largest value.

6.6. Unless the experience of a particular portfolio mirrors that of the dataset fitted using these models, the sets of μ_x produced for percentiles and sample paths may not be representative of that experience. However, in this paper we will make the assumption that we can calculate annuity values using mortality improvements based on the datasets examined.

6.7. For each percentile, we have calculated annuities using the following steps:

a) $q_{x,t}$ at exact ages were estimated from the fitted $\mu_{x,t}$ or $\mu_{x+½,t}$ as appropriate.

b) Mortality improvement rates were then estimated from these $q_{x,t}$.

c) A base table of $q'_{x,0}$ reflecting actual experience in year zero was selected and multiplied by the improvement rates to produce a two-way table of projected mortality rates by age and year $q'_{x,t}$.

d) Annuity due (ä) values payable yearly were then calculated from the $q'_{x,t}$ using a fixed 4.5% p.a. interest rate.

6.8. In these calculations we made two assumptions:

a) The highest age is 120 and $q'_{120,t} = 1$. The impact of using this assumption is not large for annuity values, except at the oldest ages, as the contribution from survival at high ages is reduced by discounting. It is more significant when calculating expectations of life.

b) The models are fitted to datasets up to age 90 (age 89 for the ONS datasets). In each annuity calculation we have assumed that mortality improvements at higher ages would be the same as at the oldest age in the dataset fitted. An alternative approach would have been to assume that improvements reduced to zero at some high age, say 120, but we have not explored the effect of this.

6.9. Our approach to calculating annuity values allows only for uncertainty relating to projected mortality improvements. We have made no allowance for stochastic variability arising from the size of the portfolio to be valued or for any heterogeneity within the portfolio. Working Paper 15 discussed the additional uncertainty arising from these sources and showed how

a) stochastic variability can be reflected in the annuity calculations by modelling the deaths as Binomial or Poisson variables based on the exposures and projected initial mortality rates and

b) heterogeneity can be reflected in the annuity calculations by separately modelling homogenous sub-groups within the portfolio.

7. P-Spline model features

Minimum data requirements

7.1. The quality and size of the dataset fitted by the model are crucial factors affecting the reliability of the fit. In our view, the data must cover a minimum period of 20 calendar years in order to provide sufficient data on which to base assumptions of longer term mortality trends. When using age-cohort penalties with P-Spline models, datasets covering longer periods may be necessary to provide sufficient numbers of cohorts with histories that are long enough to be significant.

7.2. Similarly, our view is that a minimum age-range of 40 years should be fitted in order to give the models enough data to project longer term trends. Again, when using age-cohort penalties with P-Spline models, a bigger age-range may be needed in the data.

7.3. As described in Working Paper 3, P-Spline models apply more smoothing to smaller datasets (measured by exposures and numbers of deaths) even if they cover the same periods and age ranges as a larger dataset. In order to avoid over-smoothing which may possibly ignore real trends within the data, we required a minimum number of exposures (1,000 lives) and deaths (30) in each data cell by age and year.

Impact of additional data and parameterisations

7.4. Adding data for an additional calendar year to the fitted dataset can have a significant effect on projections although the fit in the region of the data is far less affected. This feature shows how the models take account of additional information and is similar to the features observed in stochastic asset models where additional data can have a significant effect on projected returns. This is illustrated by the contour maps in Appendix B showing projected improvements, using the P-Spline model with age-cohort penalties, as additional data for the years 1983 to 2003 is added to the dataset. Appendix C shows similar contour maps for fits using the P-Spline model with age-period penalties.

7.5. The models do allow users to exercise some control on the weights applied to the latest information. The fitted coefficients for the splines in the region of the projection depend on the penalty and the coefficients of the splines that span both the data and projection regions. The best way to ensure that the projections are not unduly influenced by the most recent data is to position the knots so that no polynomial piece (i.e. the fitted curve between two internal knots) in any of the fitted splines spans both the data and projection regions.

7.6. Figure 1 uses a one-dimensional spline fit to illustrate how this can be done**.** For this purpose, splines are fitted to the period dimension and the diagram shows splines of degree 3 (i.e. cubic splines) with a non-zero value in the years 1992 to 2020 along with their internal and external knot positions. For each spline, an internal or external knot is placed every 4 years and so each spline spans across 16 years. If data is available from 1992 to 2004, the splines labelled A, B and C span across both the data and projection regions. The coefficients for the last of these three splines, C, will be based on 4 years worth of data and so no polynomial piece in any of the fitted splines spans both the data and projection regions.

Figure 1

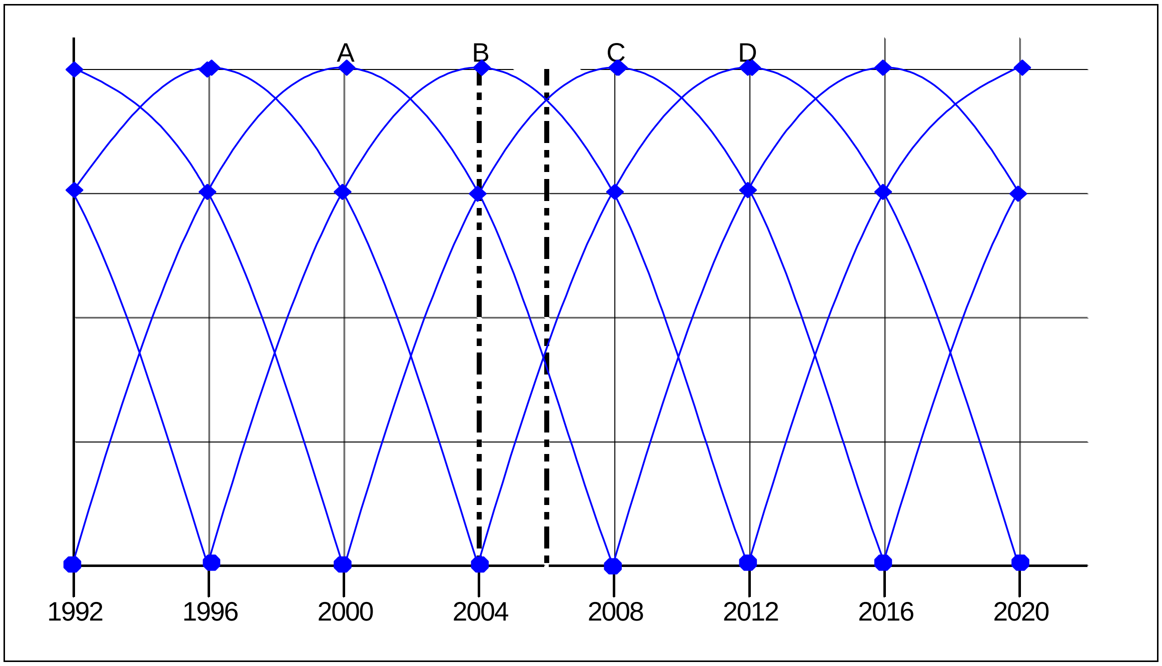

7.7. If an additional two year's worth of data is now available (i.e. 2005 and 2006) but the knot positions are unchanged, there will now be 4 splines (A to D) that span across both the data and projection regions. However, the coefficients of the last of these splines, D, will only be based on two years worth of data. Figure 2 shows that by shifting all the knot positions by 2 years so that no single polynomial piece in any of the fitted splines spans both the data and projection regions, the amount of data influencing the last spline that spans both the data and projection regions is again increased to 4 years.

7.8. Another way of considering the knot positions is in terms of minimising the number of splines that span both the data and the projection regions. In Figure 1, the introduction of additional years' data introduces a fourth spline (D). By re-positioning the knot positions as shown in Figure 2, the number of splines that span across both the data and projection regions reverts to three.

Figure 2

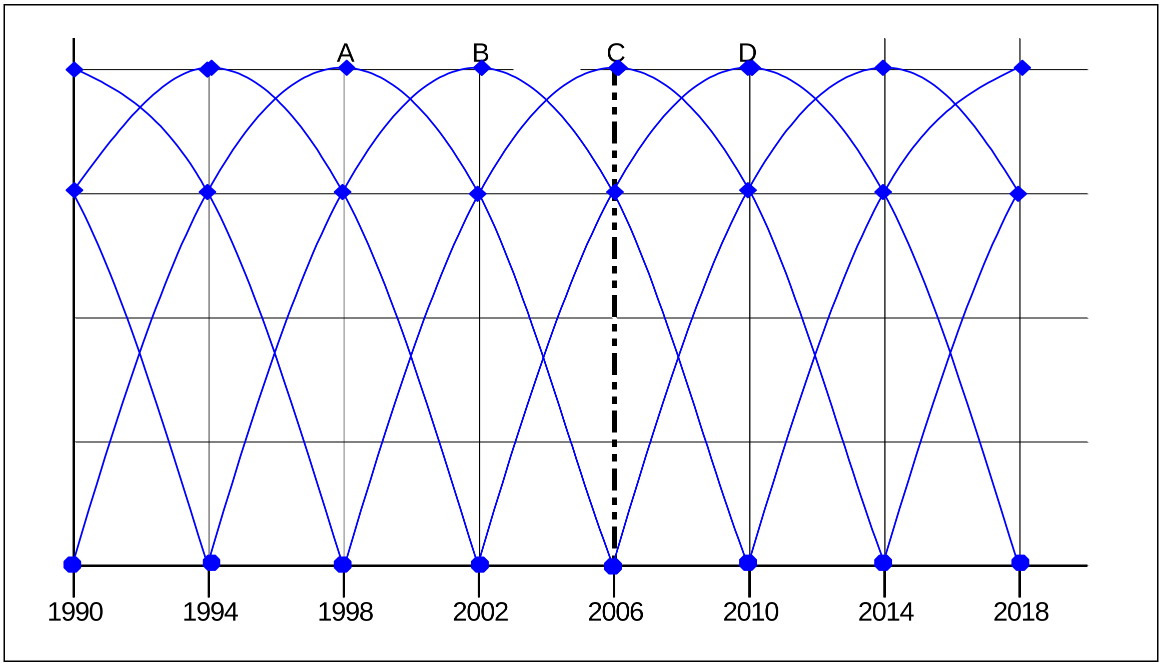

7.9. For the age-period penalties, ensuring that no single polynomial piece in any of the fitted splines spans both the data and projection regions can be achieved by positioning the knots so that a knot occurs at both the corners of the leading edge of the data (i.e. knots positioned at the highest and lowest ages in the age dimension and in the final year of the dataset in the period dimension). This may require the age range in the data to be trimmed. Figure 3 illustrates the knot positions using this approach for a dataset covering ages 60 to 76 over the years 1980 to 1992, projected to 2004.

Figure 3

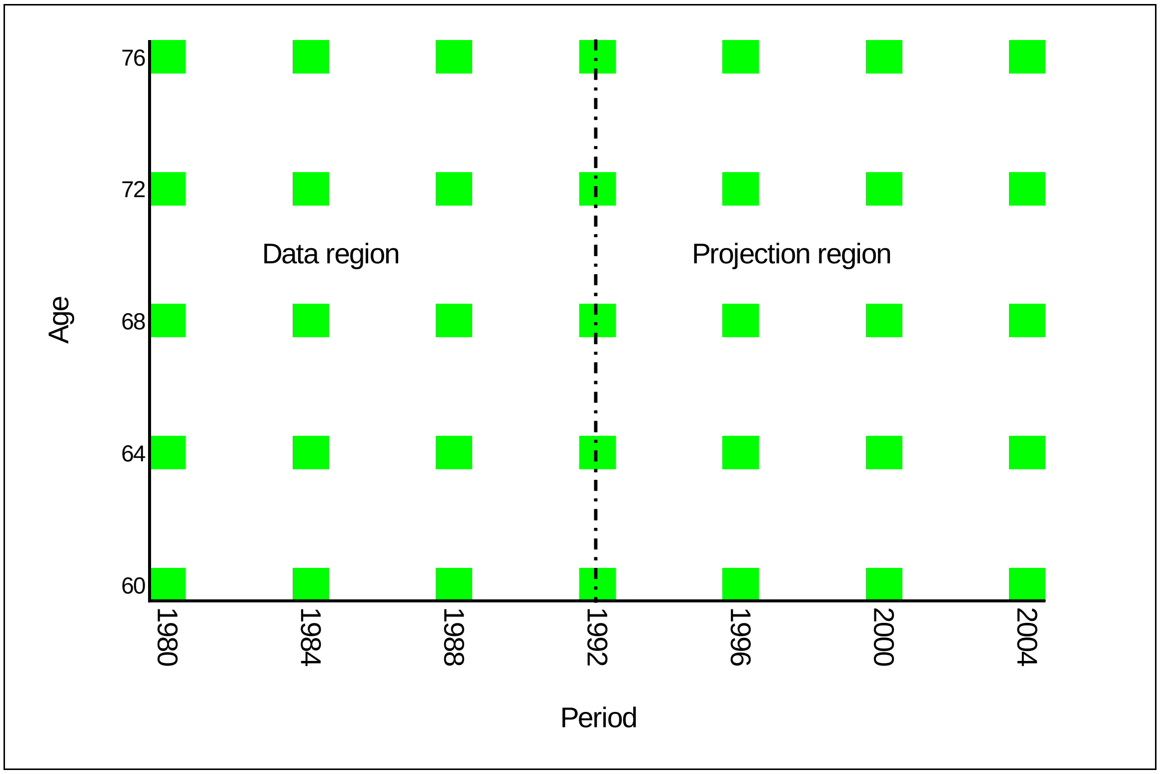

(Note: The squares in Figures 3 and 4 indicate external knots)

7.10. A similar approach can be used for the age-cohort penalties by again placing knots at both the corners of the leading edge of the data (i.e. knots at the highest and lowest ages in the age dimension and, in the cohort dimension, on the cohorts consistent with these ages in the last year of the dataset). This is illustrated in the following diagram again using a dataset covering ages 60 to 76 over the years 1980 to 1992, projected to 2004.

Figure 4

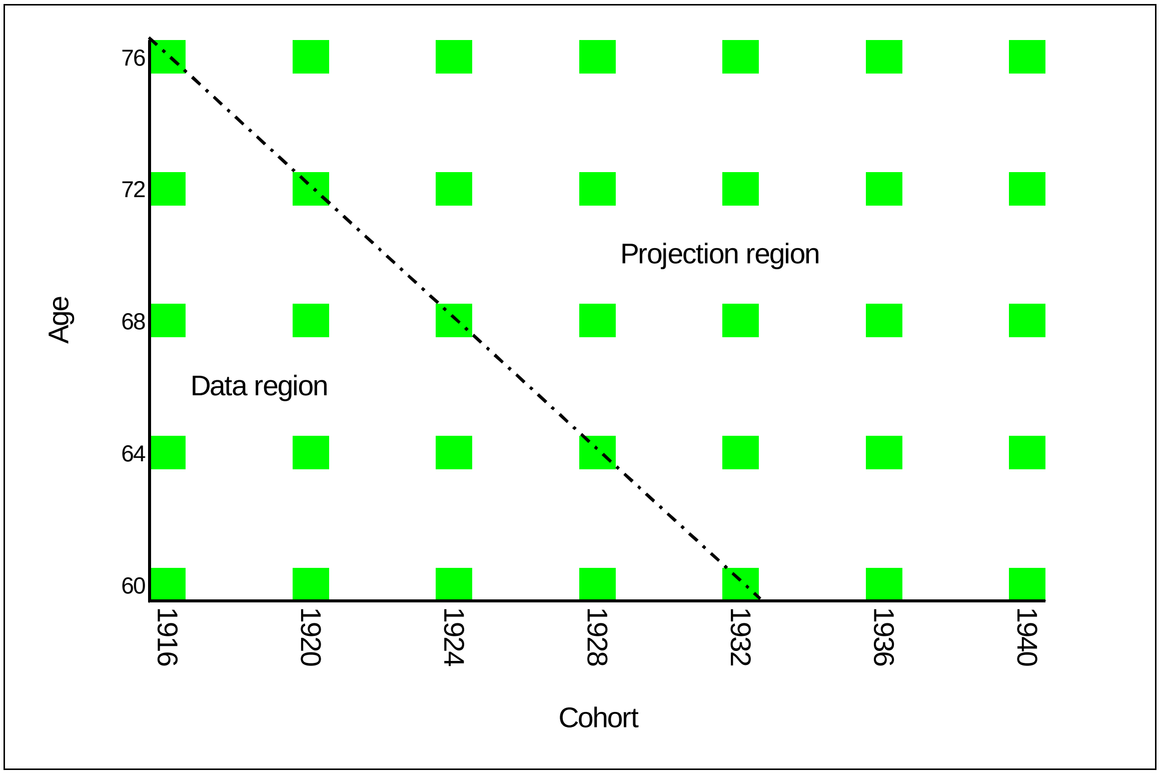

7.11. The level of smoothing applied between adjacent ages and years is driven by the spacing between knots and the penalty function. A minimal amount of spacing between knots is necessary to ensure that the coefficients of the last spline spanning both the data and projection regions are influenced by a reasonable amount of data. At the same time, knots should also not be spaced too far apart as this may smooth out recent trends in the data region and also reduce the influence of the penalty on the projection. Therefore, an element of judgement is necessary as there will be a range of choices regarding knot spacing that will give reasonable and stable projections.

7.12. However, within this range of reasonable choices for knot spacing, changing the spacing between knots will not usually materially affect the extent of smoothing in the region of the data as the penalty that will be fitted will take account of the spacing between knots. Spacing the knots closer together reduces the amount of smoothing from this source and the fitting process will fit a higher penalty value with the result that the overall level of smoothing carried out in the fit should not change significantly.

7.13. If the intention is that the splines fitted in the region of the projection should be driven primarily by the penalty, then it is important to ensure that as much as possible of the smoothing is also being carried out by the penalty rather than the knot spacing. To achieve this, within the acceptable range of knot spacing, the knots should be placed as closely together as possible. As the minimum possible spacing between knots will be driven by computational issues, a trial and error approach may be required to find the minimum knot spacing necessary for a particular dataset.

7.14. We have followed this approach with the additional requirement that the spacing between knots should not change as additional years' data is used in the projections. The parameterisations used are given in Appendix A.

7.15. Expanding the age-range in the dataset fitted can have a significant effect on projections, particularly where mortality trends in the additional age ranges differ markedly from that in adjacent age ranges. P-Spline fits are particularly affected by this feature as the fit is particularly influenced by trends at the edges of the dataset (i.e. by age and year or cohort). This is illustrated by the contour maps in Appendix D, which show projected improvements, using age-cohort penalties, as the age range of 21 to 90 is reduced to ages 42 to 90. As can be observed, the changes affect the mortality improvement rates at most ages below 80 throughout the projection.

Projecting cohort and period features

7.16. Results produced by the P-Spline model using either age-period or age-cohort penalties can identify cohort features in the region of the data as illustrated by the contour maps shown in Appendices B and C. The age-period penalties can project cohort effects if these are particularly strong in the region of the data; however fits using age-cohort penalties will more strongly project cohort effects into the future. An example of cohort features being projected by the age-cohort and age-period penalties is shown in Appendix E using the male ONS dataset to 2003.

8. What if these models had been used previously?

8.1. An obvious way to test the projection methodologies now under consideration is to look and see what would have happened if they had been used in the past, although of course this by no means provides any guarantee as to the appropriateness of any future projection. To do this we took a starting point of 1994, which was the final year of the quadrennium used to produce the "92" Series tables. That is, we have a set of graduated base tables of q_x, derived from data for 1991-94 and which is assumed to be applicable to 1992, and a dataset of deaths and exposure for each calendar year from 1947 to 1994, for the assured lives, and from 1961 to 1994 for ONS data. In this section of this report we will be using pensioner amounts base tables.

Starting in 1994

8.2. Table M1 gives some example annuity values produced from this data. To do this, P-Spline models were fitted to the data up to 1994 and projected for 1995 and years after that. In each case, a two-way table of q_x was then produced by applying improvement factors from these projections to adjusted "92" Series base tables of q_x. The base tables applying to mortality rates in 1992 were adjusted by 100A/Es for 1994 to allow for actual mortality improvements up to that year. (Note: this also has the effect of removing some of the smoothing implicit in the graduated base table). Values for annuity due ($\ddot{a}_x$) at the ages shown were then calculated as at 1995 using a 4.5% p.a. interest rate.

8.3. For comparison, the first line in the table gives annuity due values for u=1995 as calculated on the original "92" Series basis, again first adjusted by 100A/Es to 1994. Table M2 then shows the annuity values described in the previous paragraph as a percentage of these comparison numbers.

Table M1

	Projection based on male assured lives, 1947 to 1994			Projection based on male ONS data, 1961 to 1994		
	4.5% annuity value at age			4.5% annuity value at age		
Mortality basis	60	65	75	60	65	75
PMA92u95	14.388	12.549	8.686	14.388	12.549	8.686
PMA92u95p-s50ac	14.867	12.953	8.943	15.188	13.155	8.950
PMA92u95p-s97.5ac	15.187	13.230	9.113	15.908	13.760	9.281
PMA92u95p-s2.5ac	14.579	12.703	8.786	14.576	12.651	8.671
PMA92u95p-s50ap	14.818	12.917	8.931	15.543	13.604	9.372
PMA92u95p-s97.5ap	15.310	13.340	9.191	20.487	19.289	15.238
PMA92u95p-s2.5ap	14.405	12.559	8.705	10.192	9.114	6.829

Table M2

Annuity values in the shaded cells. Other cells show the values in Table M1 as percentages of values in the shaded cells.						
	Projection based on male assured lives, 1947 to 1994			Projection based on male ONS data, 1961 to 1994		
	4.5% annuity value at age			4.5% annuity value at age		
Mortality basis	60	65	75	60	65	75
PMA92u95	14.388	12.549	8.686	14.388	12.549	8.686
PMA92u95p-s50ac	103.3%	103.2%	103.0%	105.6%	104.8%	103.0%
PMA92u95p-s97.5ac	105.6%	105.4%	104.9%	110.6%	109.7%	106.9%
PMA92u95p-s2.5ac	101.3%	101.2%	101.2%	101.3%	100.8%	99.8%
PMA92u95p-s50ap	103.0%	102.9%	102.8%	108.0%	108.4%	107.9%
PMA92u95p-s97.5ap	106.4%	106.3%	105.8%	142.4%	153.7%	175.4%
PMA92u95p-s2.5ap	100.1%	100.1%	100.2%	70.8%	72.6%	78.6%

8.4. The naming convention for the mortality bases used in these tables is consistent with the conventions previously used by the CMI. The first part of the name (i.e. PMA92) refers to the base mortality table ("92" Series, pensioners, males, amounts). The following part, "u", specifies that the calculation is done using the set of q_x for lives aged 60, 65 or 70 in 1995 and following them as they age though successive calendar years to the end of the table i.e. following diagonals for particular years of birth. The next part of the basis name ("p-s") refers to the P-Spline fitting and projection method and the number appended to that is the confidence interval used in the projection. Lastly, the letters "ac" or "ap" have been added to denote that age-cohort or age-period penalties have been applied.

8.5. Looking at the results shown in the table for the assured lives dataset, the 95% confidence intervals (i.e. values between the 97.5 and 2.5 percentiles) range from 4.3% at age 60 to 3.7% at age 75 when using age-cohort penalties. When using age-period penalties, the confidence intervals range from 6.3% at age 60 to 5.6% at age 75.

8.6. The 'age-cohort' results derived from the ONS data show wider confidence intervals that range from 9.3% at age 60 to 7.1% at age 75. However the 'age-period' results from the ONS data are extremely wide, ranging from 71.6% to 96.8% between ages 60 and 75, and are clearly an indication of a badly fitting model. However, such judgments are not always obvious and the user of these methodologies will inevitably have to apply their own judgments about the results in their case. The 'age-cohort' results seem more sensible but are still a little wider than for similar assured lives data. These differences relate to differences in the data and, possibly, to the feature described in CMI Working Paper 3, whereby confidence intervals for larger datasets can be larger than for smaller ones. That is, the model gives more weight to the variability in a larger dataset than in a smaller one and thus fits more parameters in order to fit to the data more closely which results in greater parameter uncertainty.

8.7. Using age-cohort penalties, the 50th percentile annuity values calculated using projections of the male ONS data are higher than those calculated using the male assured lives data. This seems to be due to greater acceleration in mortality improvements observed in the ONS data in the 10 years to 1994 for the generation born between 1925 and 1945 which is then projected forward.

Moving to 2004

8.8. Given this approach, we can rebase the projections and calculate annuity values from 1993 onwards by adding new data for succeeding years. For example, in 1997, assuming data to 1996 is available, the P-Spline model could be refitted to that data and updated projections produced. The table of base q_x can also be adjusted to update it to 1996 by applying appropriate 100A/Es and new annuity values can then calculated. In this way the driver of the development of the mortality element of the annuity basis over successive years is the availability of the latest mortality data rather than the infrequent production of new base tables or adjustments to older tables.

8.9. Figure 5 shows 50th percentile annuity values based on a P-Spline fitting to assured lives data with an ‘age-cohort’ penalty applied. The “wiggly line” in the graph shows how an annuity value for a 60 year old male in each year would have progressed over the period 1993 to 2004 as additional data became available (i.e. the annuity values are recalculated taking account of each year of additional data). The values shown on the “wiggly line” for 1995 are those described in the table above. From the end of the “wiggly line”, in 2004, the dotted line shows the projection of future annuity due values, based on all data to 2003. The central blue line is the projected annuity values based only on the data as at 1992 (i.e. the base tables and the projections for the annuities in the line are not adjusted for actual experience in later years as it becomes available), starting the calculation in 1993. The two ‘dashed’ lines show projected annuity values based on the 1992 data but calculated from the 2.5th and 97.5th percentiles respectively. That is, they represent the confidence interval applicable to the 1992 based annuity values.

Figure 5

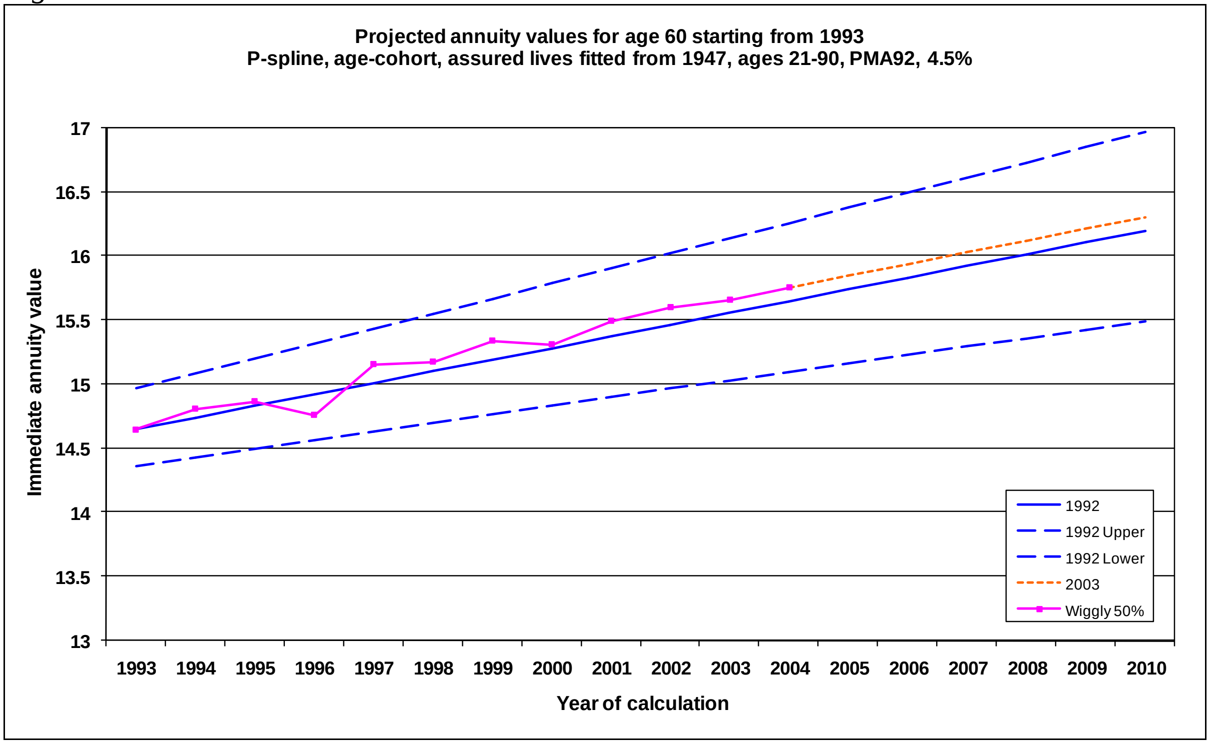

8.10. Figure 5 shows several things. A clear message is that as each year's data is added and mortality rates are projected into the future, the resulting annuity rates are rather more volatile than might have been expected. This is the case even though those projections are firmly based on, and constrained by, data from previous years. See the discussion on "knots" in paragraphs 7.4 to 7.10 for a description of how parameter settings can be used to minimise this effect. It is reassuring to see that the actual values fit comfortably within the 95% confidence interval. The Working Party has examined a number of ages to see if this result is repeated and is happy to report that it is. We found that while annuity values are volatile the direction of their future projection is not as the underlying trends in the data have not changed over this time.

8.11. Figure 6 is similar but now we have tracked the changes in the confidence intervals over time.

Figure 6

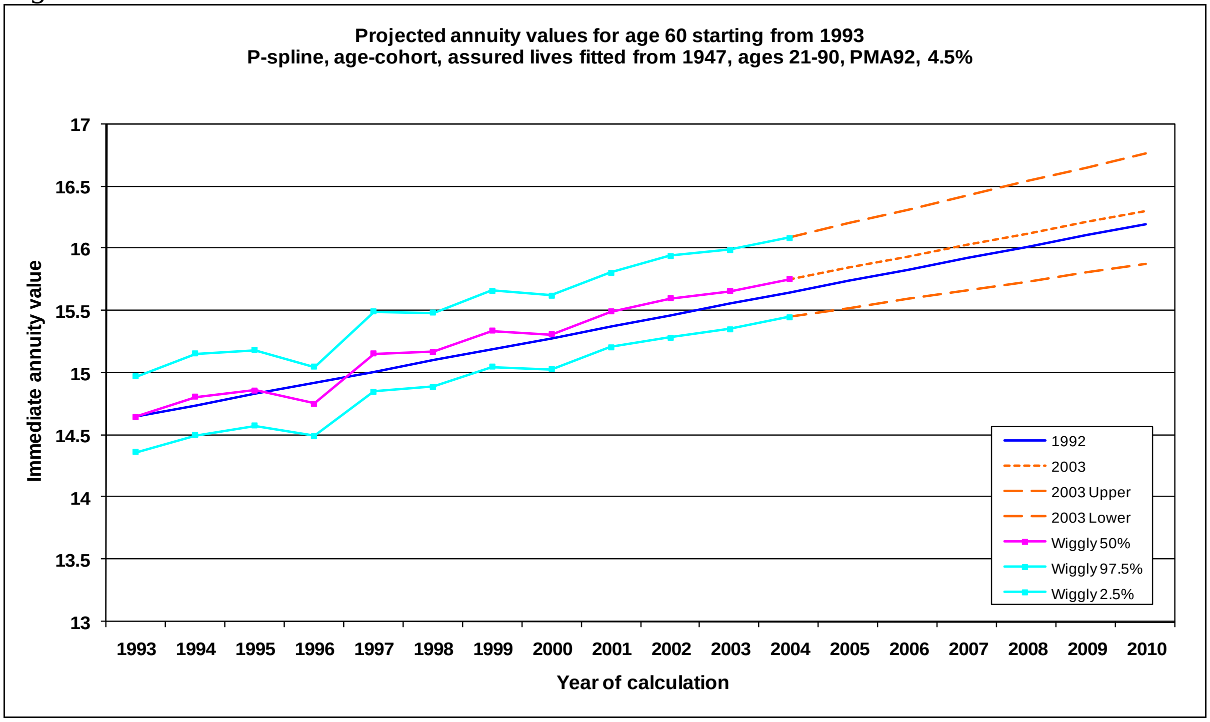

8.12. Figure 7 examines the effect of using the P-Spline model with male ONS data. The "wiggly lines" show a more volatile picture reflecting the more rapid mortality improvements observed in the ONS data between 1996 and 2002 compared to the blue line based on projections using ONS data to 1992 only.

Figure 7

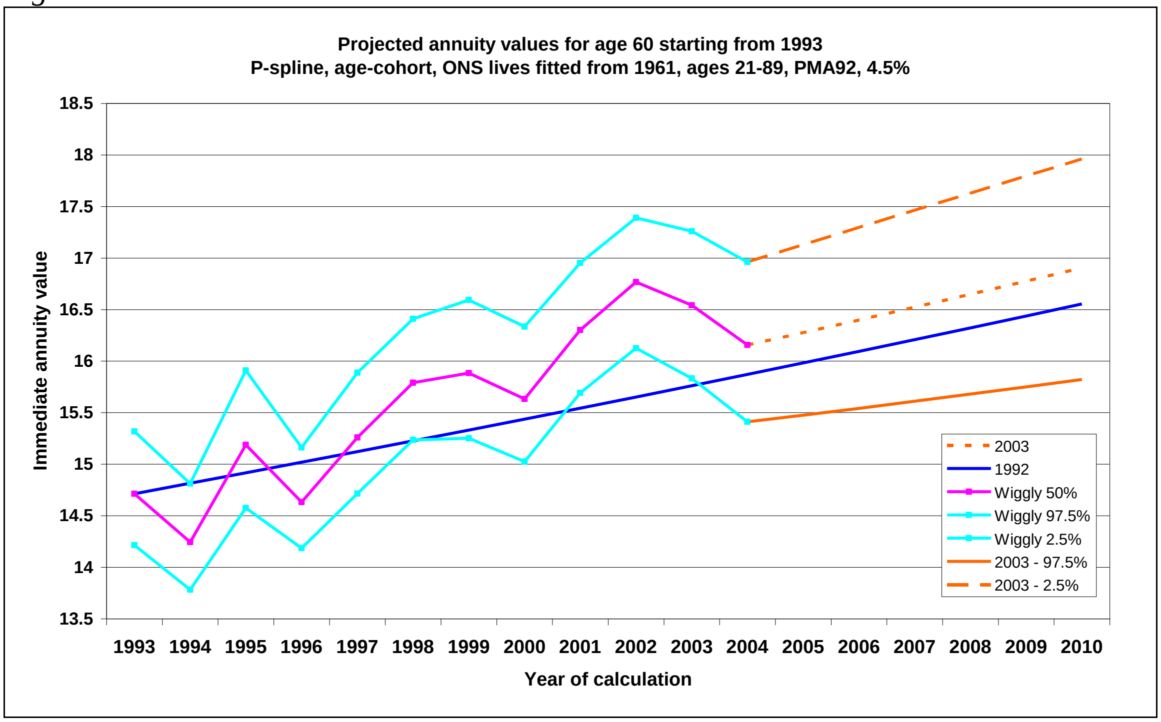

8.13. We can also examine the effect of these model fits on annuity rates over longer periods. Figure 8 shows similar information for annuity values over the period from 1984 to 2004. In this case the age used is 75. The wiggly line shows the progression in $\ddot{a}_{75}$ over the period, the lines are the 50th percentile projections based on data to 1983, 1992 and 2003 respectively. This graph shows that, when 1992 data becomes available in 1993, the underlying trends projected by the model change as the model starts to project forward the cohort effect that is becoming apparent in the data.

Figure 8

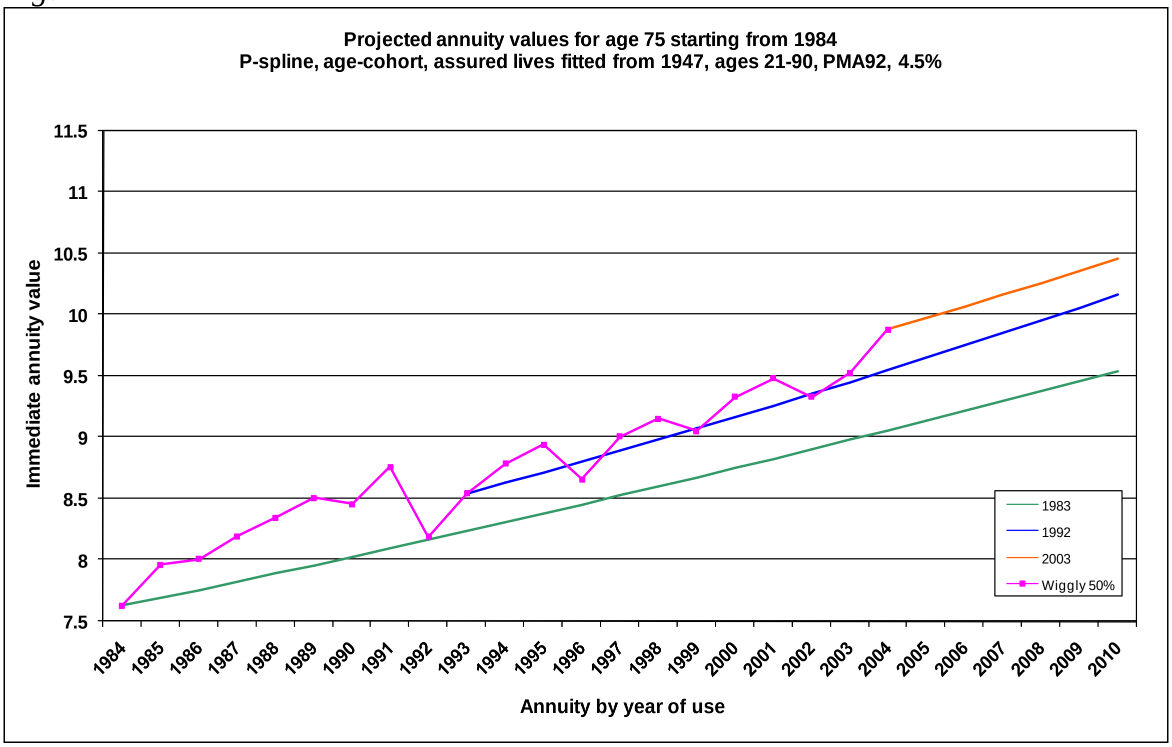

The position in 2004

8.14. Table M3 shows the position in 2004. The first line of figures shows annuity values derived by applying the medium cohort projection published in November 2002 to the "92" Series base table. This is included for comparison purposes as these are the only figures in this table previously publicly available. The next three rows show annuity values again calculated using the "92" Series base table but now updated to 2003 by applying actual improvement factors (100A/Es) applicable for 2003 then with the projected improvements using the long, medium and short cohort projections applied from 2004.

8.15. Subsequent rows of Table M3 then show annuity values for 2004 based on male assured lives and ONS data as at 2003, calculated using the method described in paragraph 8.2. As these projections are based on data to 2003, they are not strictly comparable with the annuity values described in the previous paragraph as they were calculated using projections based on data only to 2000. However, for the purpose of this paper this difference is not material. Table M4 shows the annuity values described as a percentage of the annuity values calculated on the "92" Series basis with medium cohort adjustment and adjusted for experience to 2003.

Table M3

	Projection based on male assured lives, 1947 to 2003			Projection based on male ONS data, 1961 to 2003		
	4.5% annuity value at age			4.5% annuity value at age		
Mortality basis	60	65	75	60	65	75
PMA92u04mc (unadjusted)	15.480	13.786	9.842	15.480	13.786	9.842
PMA92u04mc	15.218	13.640	9.902	15.218	13.640	9.902
PMA92u04lc	15.620	14.154	10.355	15.620	14.154	10.355
PMA92u04sc	15.044	13.415	9.599	15.044	13.415	9.599
PMA92u04p-s50ac	15.756	14.012	9.880	16.156	14.409	10.029
PMA92u04p-s97.5ac	16.088	14.309	10.073	16.961	15.151	10.484
PMA92u04p-s2.5ac	15.452	13.742	9.703	15.411	13.763	9.651
PMA92u04p-s50ap	15.746	14.027	9.913	14.599	13.020	9.174
PMA92u04p-s97.5ap	16.277	14.503	10.222	16.359	14.508	9.989
PMA92u04p-s2.5ap	15.291	13.621	9.645	13.653	12.203	8.676

Table M4

Annuity values in the shaded cells. Other cells show the values in Table M3 as percentages of values in the shaded cells.						
	Projection based on male assured lives, 1947 to 2003			Projection based on male ONS data, 1961 to 2003		
	4.5% annuity value at age			4.5% annuity value at age		
Mortality basis	60	65	75	60	65	75
PMA92u04mc	15.218	13.640	9.902	15.218	13.640	9.902
PMA92u04lc	102.6%	103.8%	104.6%	102.6%	103.8%	104.6%
PMA92u04sc	98.9%	98.3%	96.9%	98.9%	98.3%	96.9%
PMA92u04p-s50ac	103.5%	102.7%	99.8%	106.2%	105.6%	101.3%
PMA92u04p-s97.5ac	105.7%	104.9%	101.7%	111.5%	111.1%	105.9%
PMA92u04p-s2.5ac	101.5%	100.7%	98.0%	101.3%	100.9%	97.5%
PMA9204up-s50ap	103.5%	102.8%	100.1%	95.9%	95.4%	92.6%
PMA92u04p-s97.5ap	107.0%	106.3%	103.2%	107.5%	106.4%	100.9%
PMA92u04p-s2.5ap	100.5%	99.9%	97.4%	89.7%	89.5%	87.6%

8.16. Using age-cohort penalties, the annuity values calculated using projections of the male ONS data are again generally higher than those calculated using the male assured lives data. Similar to the position in 1994 as described in paragraph 8.7, this seems to be due to the higher observed mortality improvements in the ONS data over the 20 years to 2003 for the generation born between 1925 and 1945. However, this time rapid improvements for the younger generations born between 1964 and 1974 also exist. (These can be observed by comparing Figure E1 from Appendix E with Figure B11, from Appendix B).

8.17. The confidence intervals for the projections based on the assured lives dataset are little changed from the position in 1995 shown in Table M2. However, using the ONS dataset, the confidence intervals using age-period penalties are much narrower, indicating a better fitting model. Conversely the confidence intervals using age-cohort penalties for the ONS data are a little wider compared to the position in 1995, though still much narrower than using age-period penalties.

8.18. The last step, to bring the sequence completely up to date, is to look at the effect of changing the base table, from the "92" Series to the draft "00" Series [see CMI (2005d)]. To do this the next table shows "00" Series annuity values expressed as a percentage of the updated "92" Series values given in Table M3.

As might be expected, given the application of 100A/Es to the “92” Series base tables, this table shows little difference between these annuity values.

Table M5

Annuity values @ 4.5% - “00” Series as a percentage of “92” Series, base tables adjusted to 2003						
	Projection based on male assured lives, 1947 to 2003			Projection based on male ONS data, 1961 to 2003		
	4.5% annuity value at age			4.5% annuity value at age		
Mortality basis	60	65	75	60	65	75
PMA00u04p-s50ac	100.0%	100.0%	99.8%	100.0%	100.0%	99.7%
PMA00u04p-s97.5ac	100.0%	100.0%	99.7%	100.0%	99.9%	99.7%
PMA00u04p-s2.5ac	100.0%	100.0%	99.8%	100.0%	100.0%	99.8%
PMA00u04p-s50ap	100.0%	100.0%	99.8%	100.1%	100.1%	99.8%
PMA00u04p-s97.5ap	100.0%	100.0%	99.7%	100.0%	99.9%	99.7%
PMA00u04p-s2.5ap	100.0%	100.0%	99.8%	100.1%	100.1%	99.9%

The position after 2004

8.19. The next three graphs (Figures 9 to 11) show a comparison of annuity values calculated using the interim cohort projections (from 2000 onwards) and using the P-Spline projections, for males at ages 60, 65 and 75. In each case, the projections have been updated using 100 A/Es for each year up to 2004. The “wiggly” lines are as per those in the graphs in Sections 8.11 and 8.13 for ages 60 and 75 respectively.

Figure 9

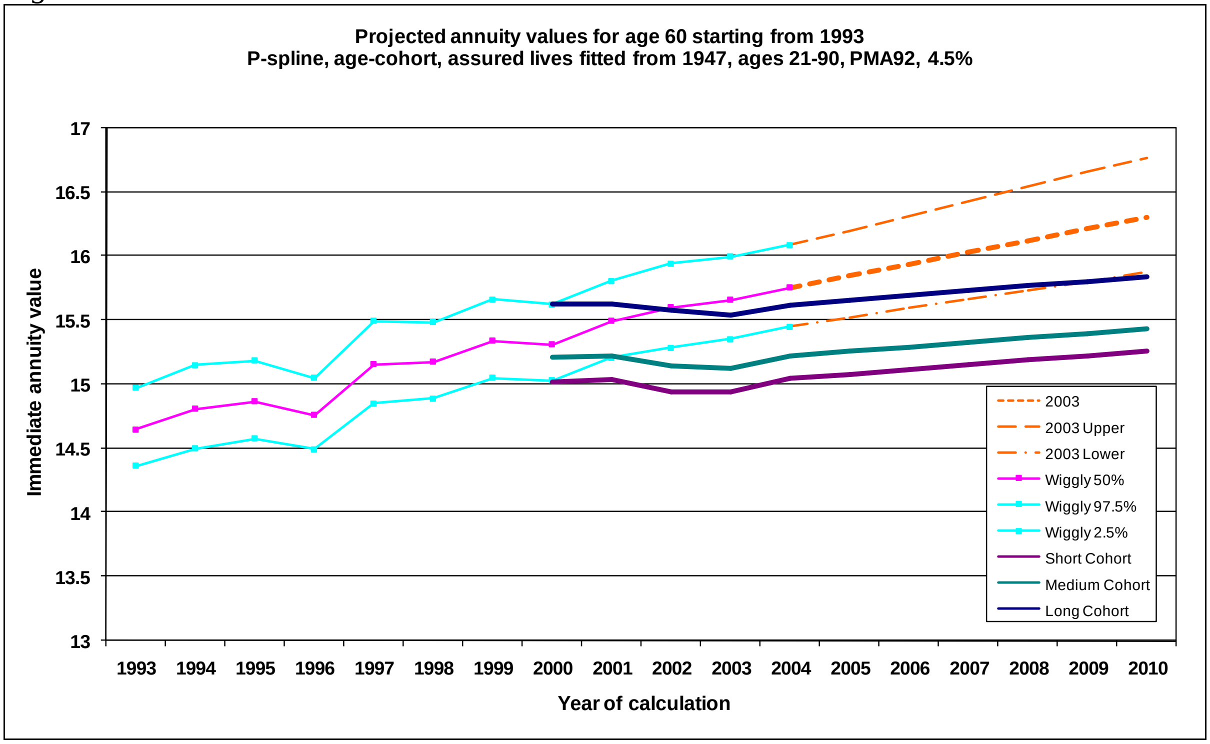

Figure 10

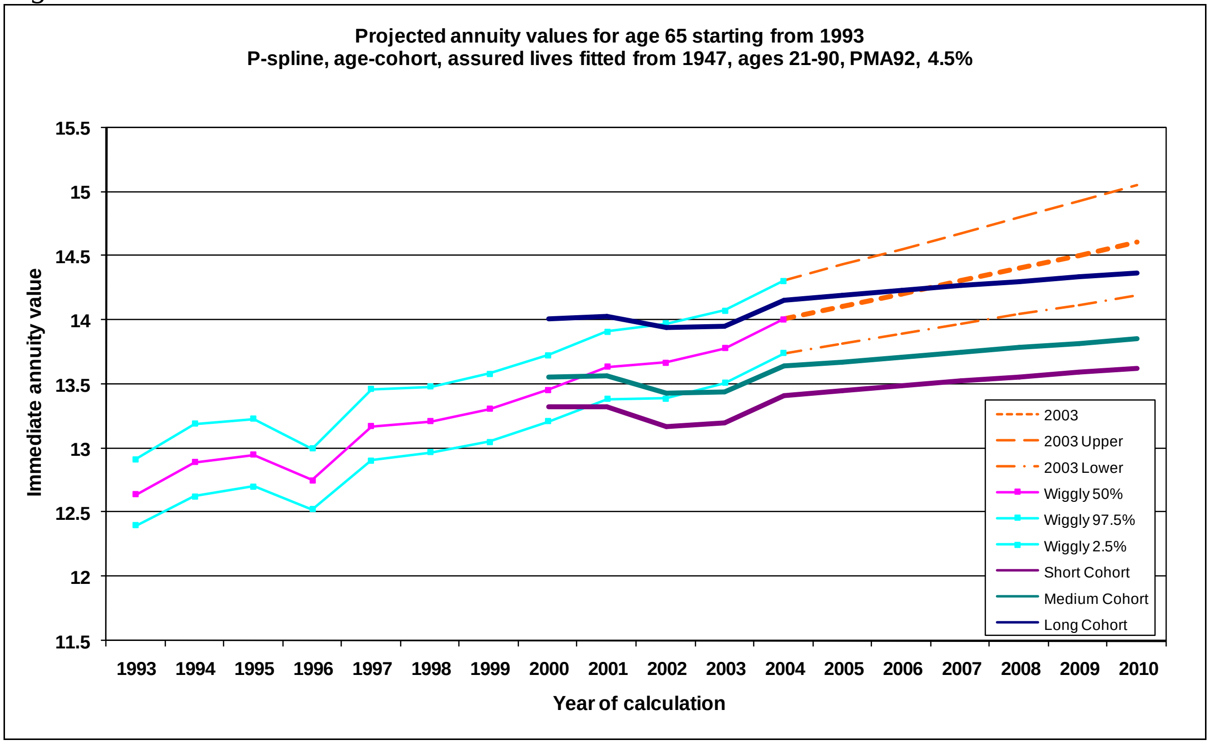

Figure 11

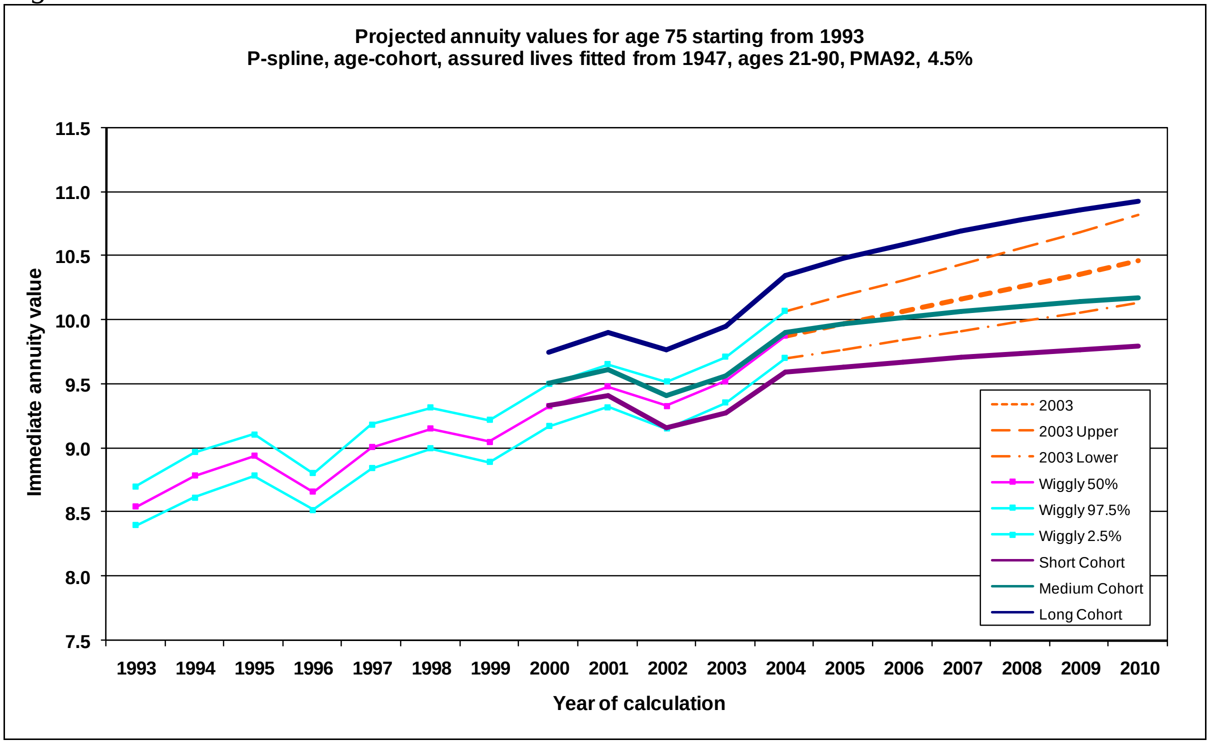

8.20. At age 75, the confidence interval indicated by the long and short cohort projections is wider than the 95% confidence interval calculated using the P-Spline projections. The medium cohort projections, though, give lower annuity values than the 50th percentile P-Spline projection from 2006 onwards. At younger ages, even the long cohort projections give lower annuity values than the 50th percentile P-Spline projection. This happens from 2003 at age 60 and from 2008 at age 65. At age 60, after 2010, the long cohort projections give annuity values that are lower than the 2.5th percentile P-Spline projection.

8.21. The interim cohort projections are more out of line with the P-Spline projections for younger generations. This is to be expected as the interim cohort projections only allowed for the high mortality improvements for the generation born around 1926 while the P-Spline projections also allow for the cohort effects of later generations based on the actual historic experience observed.

Females

8.22. Tables F1 to F5 below show annuity values using “92” Series female base mortality but projections based on male assured lives and female ONS data, similar to the male results in Tables M1 to M5. As explained in paragraph 5.6, we believe that there is insufficient female assured lives data to carry out reliable P-Spline projections.

Table F1

Mortality basis	Projection based on male assured lives, 1947 to 1994			Projection based on female ONS data, 1961 to 1994		
	4.5% annuity value at age			4.5% annuity value at age		
	60	65	75	60	65	75
PFA92u95	15.251	13.700	9.969	15.251	13.700	9.969
PFA92u95p-s50ac	15.819	14.205	10.317	15.873	14.236	10.297
PFA92u95p-s97.5ac	16.170	14.526	10.539	16.369	14.677	10.574
PFA92u95p-s2.5ac	15.496	13.911	10.114	15.422	13.840	10.049
PFA92u95p-s50ap	15.771	14.169	10.303	15.802	14.306	10.557
PFA92u95p-s97.5ap	16.315	14.666	10.646	20.393	19.522	15.972
PFA92u95p-s2.5ap	15.299	13.741	10.005	10.869	9.970	7.703

Table F2

Annuity values in the shaded cells. Other cells show the values from Table F1 as percentages of values in the shaded cells.

Mortality basis	Projection based on male assured lives, 1947 to 1994			Projection based on female ONS data, 1961 to 1994		
	4.5% annuity value at age			4.5% annuity value at age		
	60	65	75	60	65	75
PFA92u95	15.251	13.700	9.969	15.251	13.700	9.969
PFA92u95p-s50ac	103.7%	103.7%	103.5%	104.1%	103.9%	103.3%
PFA92u95p-s97.5ac	106.0%	106.0%	105.7%	107.3%	107.1%	106.1%
PFA92u95p-s2.5ac	101.6%	101.5%	101.5%	101.1%	101.0%	100.8%
PFA92u95p-s50ap	103.4%	103.4%	103.4%	103.6%	104.4%	105.9%
PFA92u95p-s97.5ap	107.0%	107.1%	106.8%	133.7%	142.5%	160.2%
PFA92u95p-s2.5ap	100.3%	100.3%	100.4%	71.3%	72.8%	77.3%

8.23. In 1995, the annuity values calculated using P-Spline projections on female ONS data are generally lower percentages of annuity values calculated using the “92” Series projections when compared to the equivalent percentages for male annuity values in Table M2. Similar to Table M2, the confidence intervals shown in Table F2 for the projections using female ONS data to 1994 and using age-period penalties are extremely wide, ranging from 62.4% to 82.9% between ages 60 and 75, clearly an indication of a badly fitting model.

8.24. Similar to the case with male ONS data, once female ONS data to 2003 is available, Table F4 shows that the confidence intervals narrow considerably for age-period penalties and, using age-cohort penalties, are comparable to the projections based on male assured lives. Again, the annuity values calculated using P-Spline projections on female ONS data are consistently lower percentages of annuity values calculated using the interim cohort projections when compared to the equivalent percentages for male annuity values in Table M4. This seems to be due to the less rapid acceleration in female mortality improvements in the 10 years to 2003, compared to the male experience, which are then projected forward.

Table F3

	Projection based on male assured lives, 1947 to 2003			Projection based on female ONS data, 1961 to 2003		
	4.5% annuity value at age			4.5% annuity value at age		
Mortality basis	60	65	75	60	65	75
PFA92u04mc (unadjusted)	16.327	14.814	11.166	16.327	14.814	11.166
PFA92u04mc	15.976	14.482	10.819	15.976	14.482	10.819
PFA92u04lc	16.400	15.020	11.333	16.400	15.020	11.333
PFA92u04sc	15.804	14.260	10.515	15.804	14.260	10.515
PFA92u04p-s50ac	16.592	14.947	10.898	16.515	14.906	10.767
PFA92u04p-s97.5ac	16.942	15.273	11.130	16.962	15.311	11.027
PFA92u04p-s2.5ac	16.265	14.645	10.685	16.101	14.536	10.531
PFA92u04p-s50ap	16.589	14.967	10.937	14.776	13.247	9.563
PFA92u04p-s97.5ap	17.147	15.490	11.308	15.753	14.085	10.067
PFA92u04p-s2.5ap	16.090	14.509	10.611	14.122	12.673	9.193

Table F4

Annuity values in the shaded cells. Other cells show the values in Table F3 as percentages of values in the shaded cells.						
	Projection based on male assured lives, 1947 to 2003			Projection based on female ONS data, 1961 to 2003		
	4.5% annuity value at age			4.5% annuity value at age		
Mortality basis	60	65	75	60	65	75
PFA92u04mc	15.976	14.482	10.819	15.976	14.482	10.819
PFA92u04lc	102.7%	103.7%	104.7%	102.7%	103.7%	104.7%
PFA92u04sc	98.9%	98.5%	97.2%	98.9%	98.5%	97.2%
PFA92u04p-s50ac	103.9%	103.2%	100.7%	103.4%	102.9%	99.5%
PFA92u04p-s97.5ac	106.0%	105.5%	102.9%	106.2%	105.7%	101.9%
PFA92u04p-s2.5ac	101.8%	101.1%	98.8%	100.8%	100.4%	97.3%
PFA92u04p-s50ap	103.8%	103.4%	101.1%	92.5%	91.5%	88.4%
PFA92u04p-s97.5ap	107.3%	107.0%	104.5%	98.6%	97.3%	93.0%
PFA92u04p-s2.5ap	100.7%	100.2%	98.1%	88.4%	87.5%	85.0%

Table F5

Annuity values @ 4.5% - "00" Series as a percentage of "92" Series, base tables adjusted to 2003						
	Projection based on male assured lives, 1947 to 2003			Projection based on female ONS data, 1961 to 2003		
	4.5% annuity value at age			4.5% annuity value at age		
Mortality basis	60	65	75	60	65	75
PFA00u04p-s50ac	99.5%	99.8%	99.5%	99.5%	99.8%	99.5%
PFA00u04p-s97.5ac	99.5%	99.8%	99.5%	99.5%	99.8%	99.5%
PFA00u04p-s2.5ac	99.5%	99.8%	99.6%	99.5%	99.8%	99.6%
PFA00u04p-s50ap	99.5%	99.8%	99.5%	99.6%	99.9%	99.7%
PFA00u04p-s97.5ap	99.5%	99.8%	99.5%	99.5%	99.8%	99.6%
PFA00u04p-s2.5ap	99.5%	99.8%	99.6%	99.6%	99.9%	99.7%

9. Conclusion

9.1. The Working Party hopes that this paper provides practical help to others seeking to experiment with P-Spline mortality projections.

9.2. The back-testing described in Section 8 demonstrates that the projections would have worked well in recent years. This of course provides no guarantee for the future.

9.3. Whilst we feel they provide a useful tool and a data-driven basis for future projections, users should nevertheless recognise the discretion involved, in the fitting parameters and particularly in the choice of model. In our work on CMI data and the England & Wales population data, cohort effects exist even when using an age-period method, which suggests that these are genuine features. However in general, the age-cohort model will project cohort effects more strongly than the age-period model and the choice is a judgment call.

9.4. The pattern of higher observed mortality improvements for P-Spline projections compared to the interim cohort projections means that the interim cohort projections based on data to 2000 are unlikely to be suitable as more recent data becomes available. The P-Spline methodology is better able to project forward the actual improvements as more recent data becomes available. It is a powerful tool for modelling mortality improvements and in aiding understanding of the associated uncertainty. Clearly, great care is needed in the choice regarding the dataset to project as well as the parameters and penalties used.

9.5. Much further work is required but that is beyond the scope of the Working Party. This further work, described in Working Paper 3, would include consideration of questions on model uncertainty, the correlation between mortality and investment risk and moving the projection methodology towards cause-specific projections. However, better data on cause-specific mortality may be required before there can be much further progress on cause-specific projections.

10. Appendices

Appendix A - Parameters used to generate projections

We have used b-splines of degree 3 and penalty order of 2 for all our fits.

Age-Cohort model

For datasets fitted using the age-cohort model the following parameters were used:

	Assured Lives Males	**ONS Males**	**ONS Females**
Age range	21-90	21-89	24-89
Knot spacing:			
- age dimension	Every 3 years	Every 4 years	Every 5 years
- cohort dimension	Every 3 years	Every 4 years	Every 5 years
Fixed knot positions:			
- age dimension	90	89	89
- cohort dimension	Last year of data less 90	Last year of data less 89	Last year of data less 89
Minimum for penalty:			
- age dimension	0.0001	0.0001	0.0001
- cohort dimension	0.0001	0.0001	0.0001
Starting point for penalty:			
- age dimension	100	100	100
- cohort dimension	100	100	100

Age-period model

For datasets fitted using the age-period model the following parameters were used:

	Assured Lives Males	**ONS Males**	**ONS Females**
Age range	22-90	23-89	23-89
Knot spacing:			
- age dimension	Every 4 years	Every 6 years	Every 6 years
- period dimension	Every 4 years	Every 6 years	Every 6 years
Fixed knot positions:			
- age dimension	90	89	89
- period dimension	Last year of data	Last year of data	Last year of data
Minimum for penalty:			
- age dimension	0.0001	0.0001	0.0001
- period dimension	0.0001	0.0001	0.0001
Starting point for penalty:			
- age dimension	100	100	100
- period dimension	100	100	100

Appendix B – Contour maps of projected improvements for male assured lives using age-cohort penalties

Figure B1

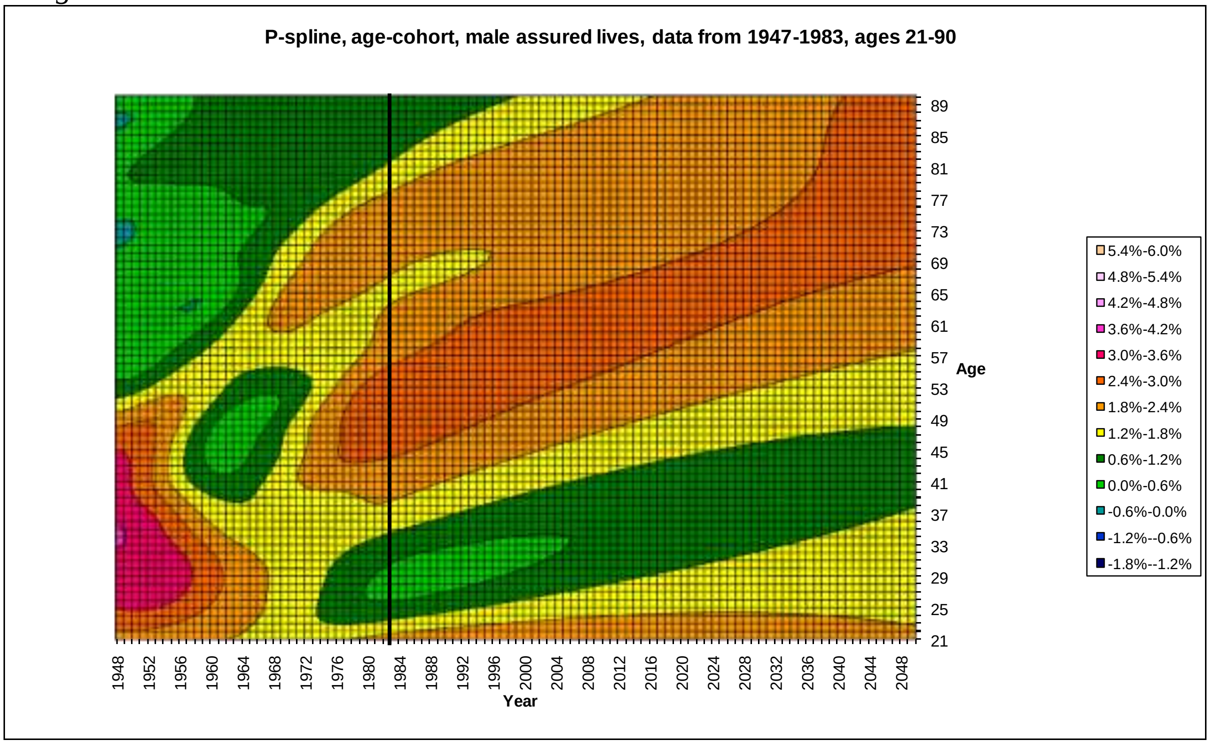

Figure B2

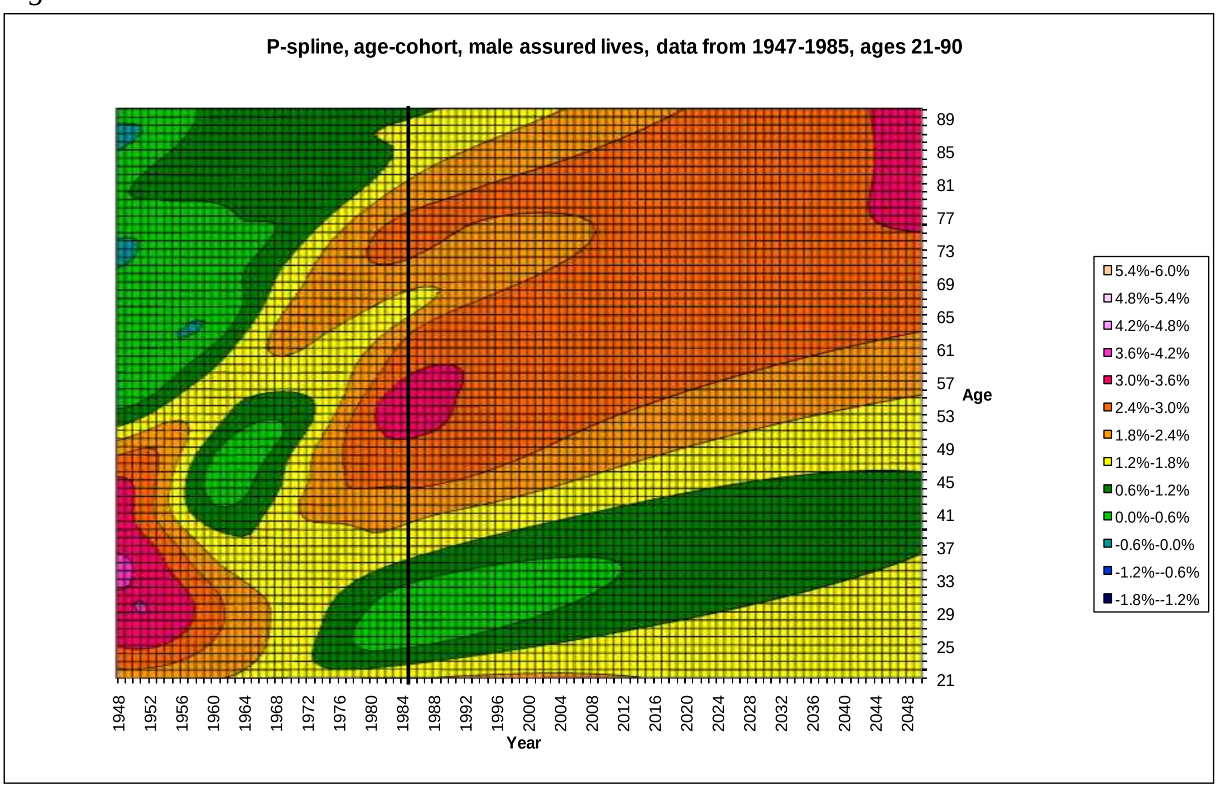

Figure B3

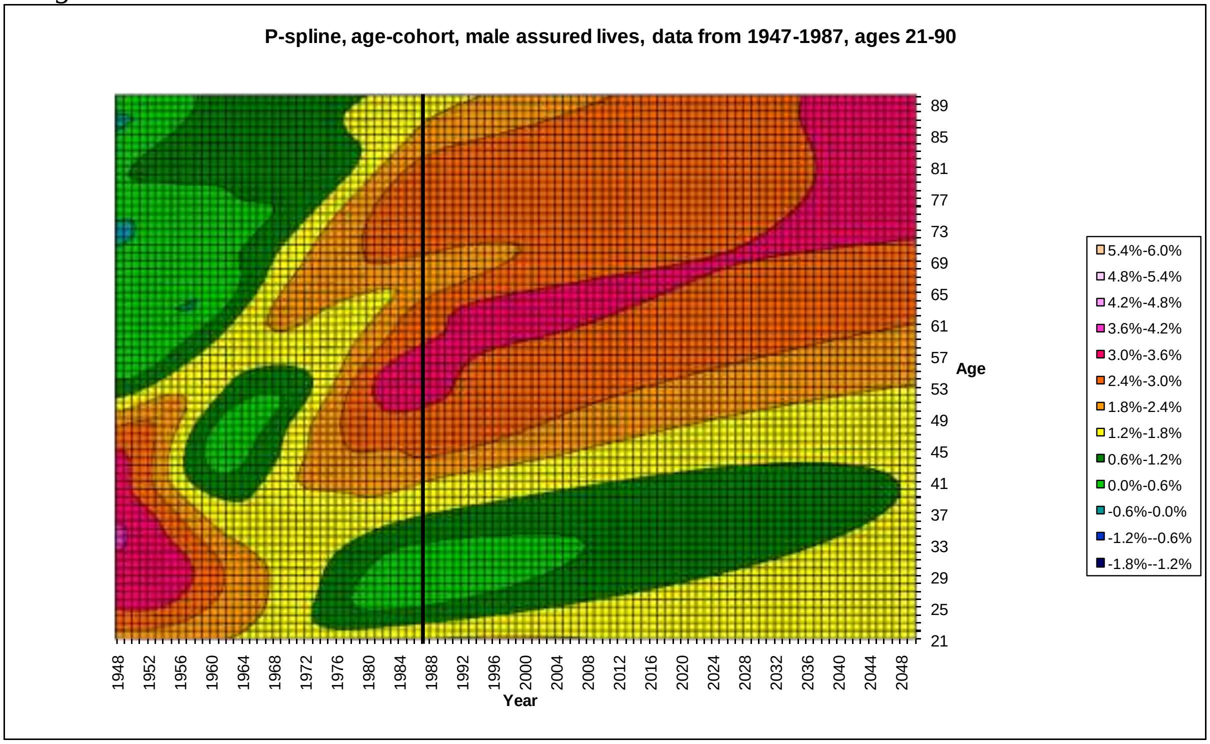

Figure B4

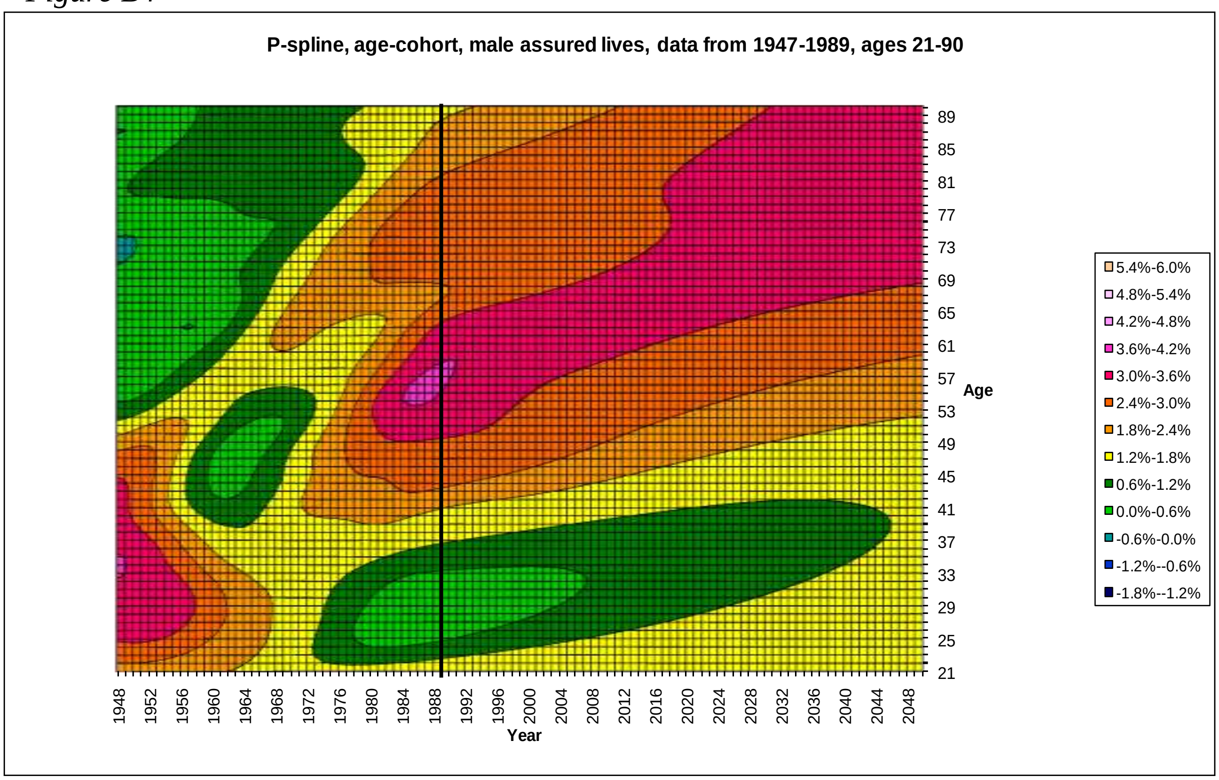

Figure B5

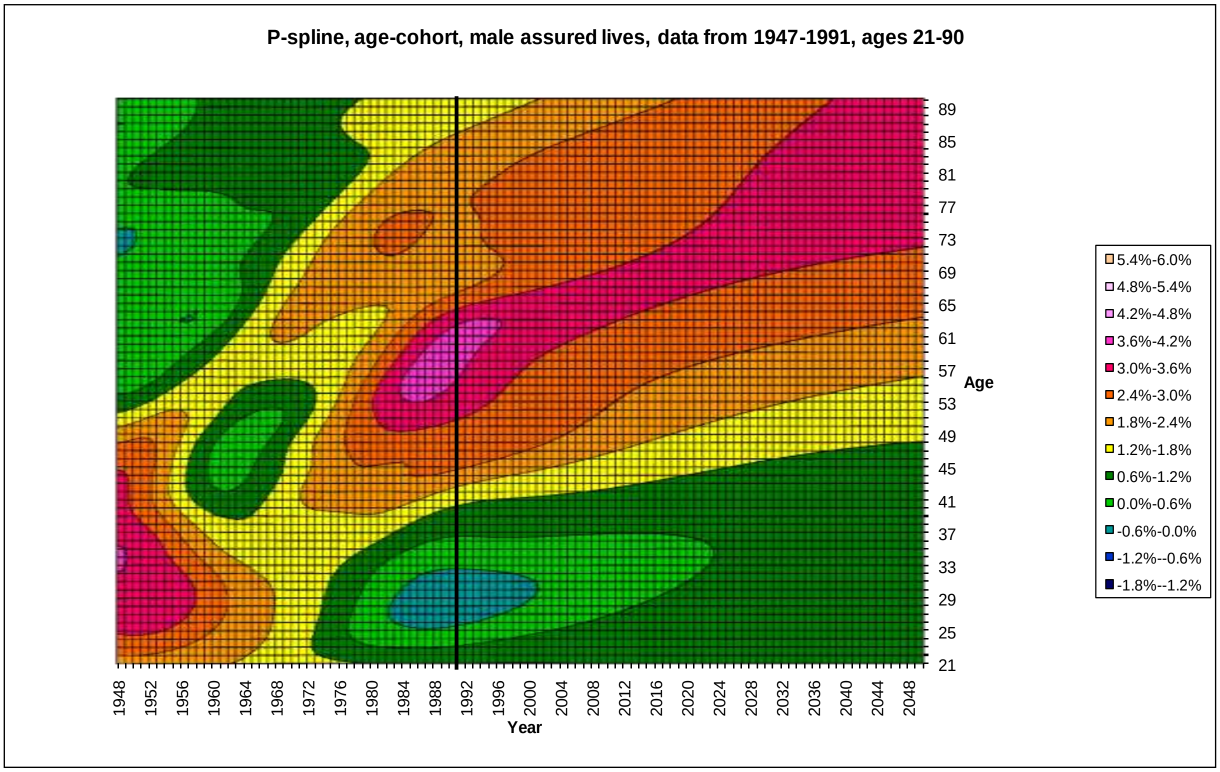

Figure B6

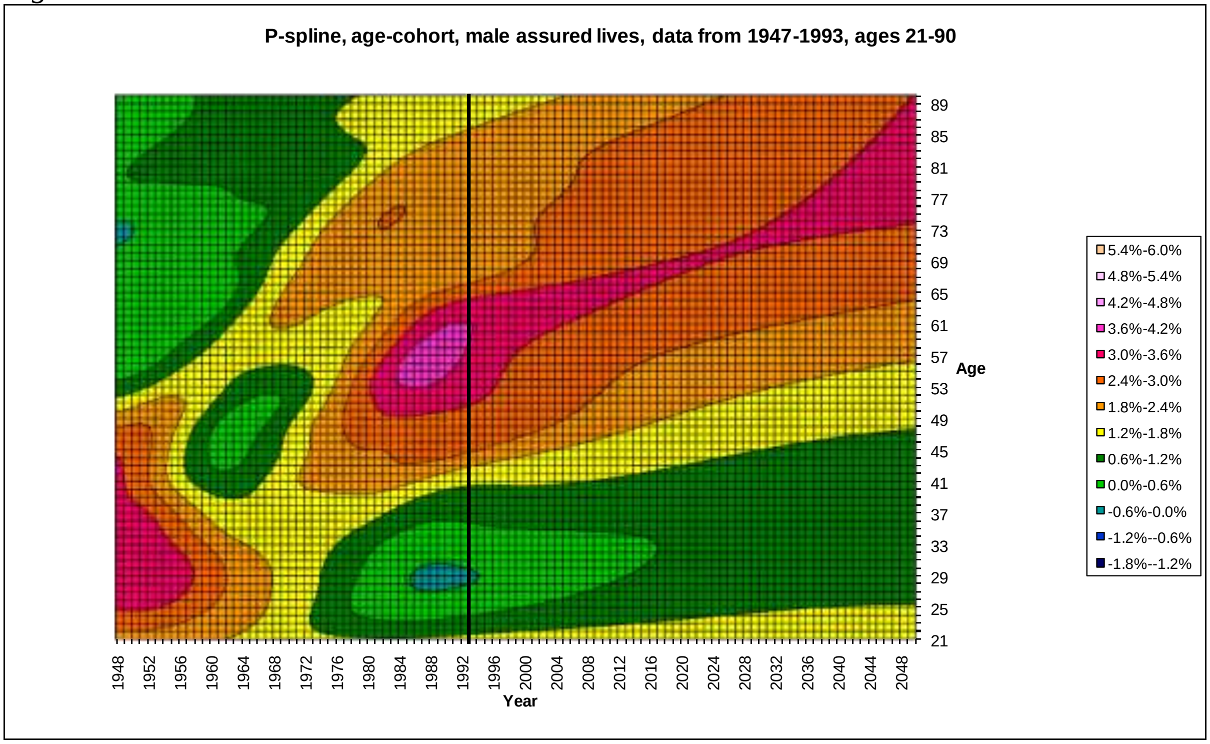

Figure B7

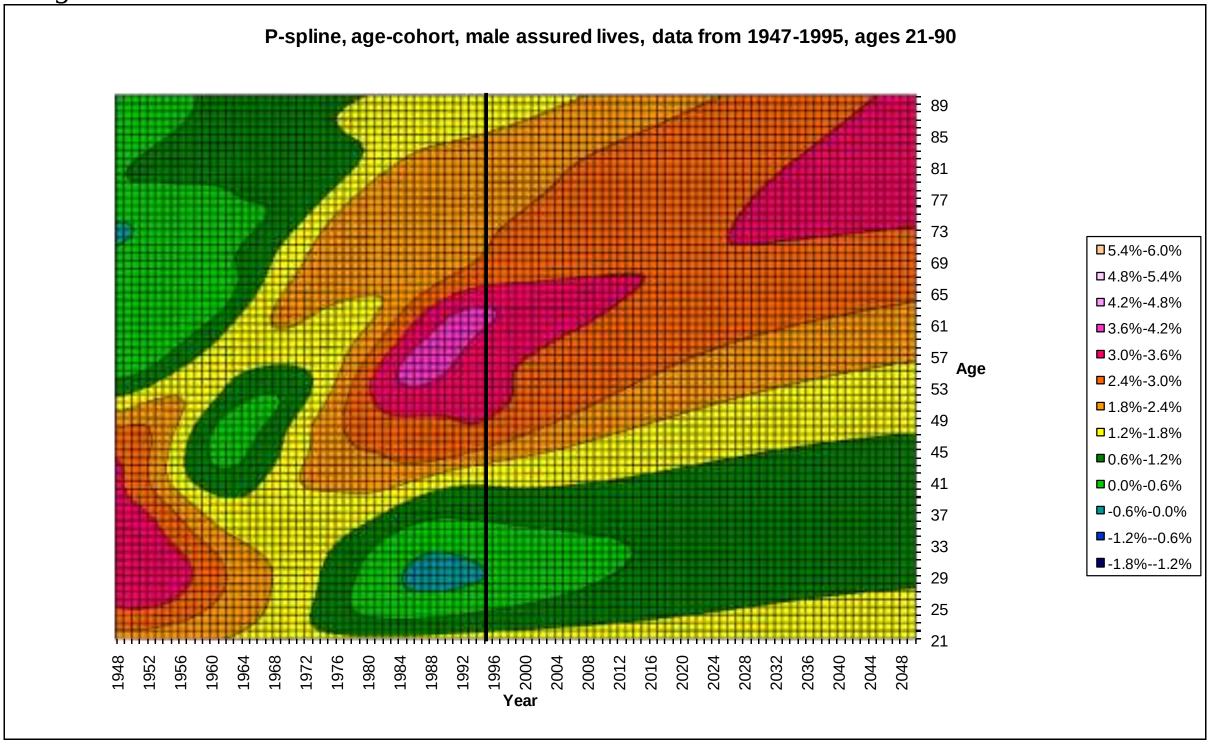

Figure B8

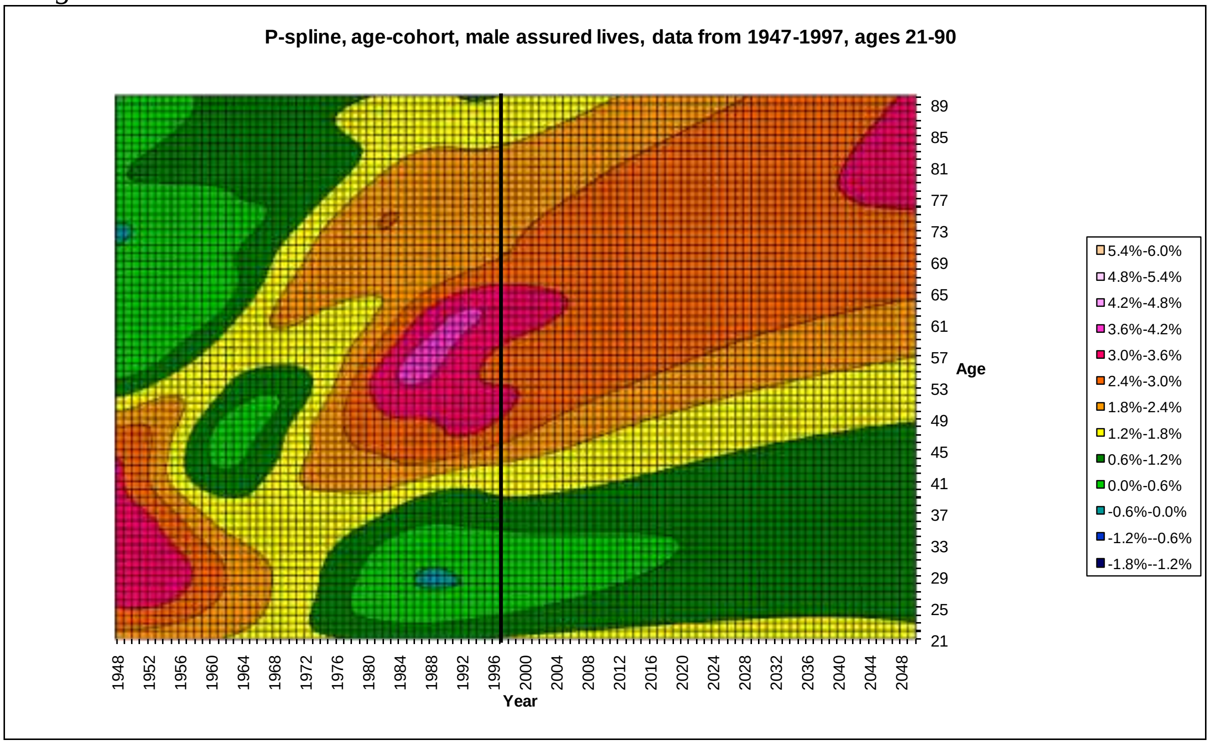

Figure B9

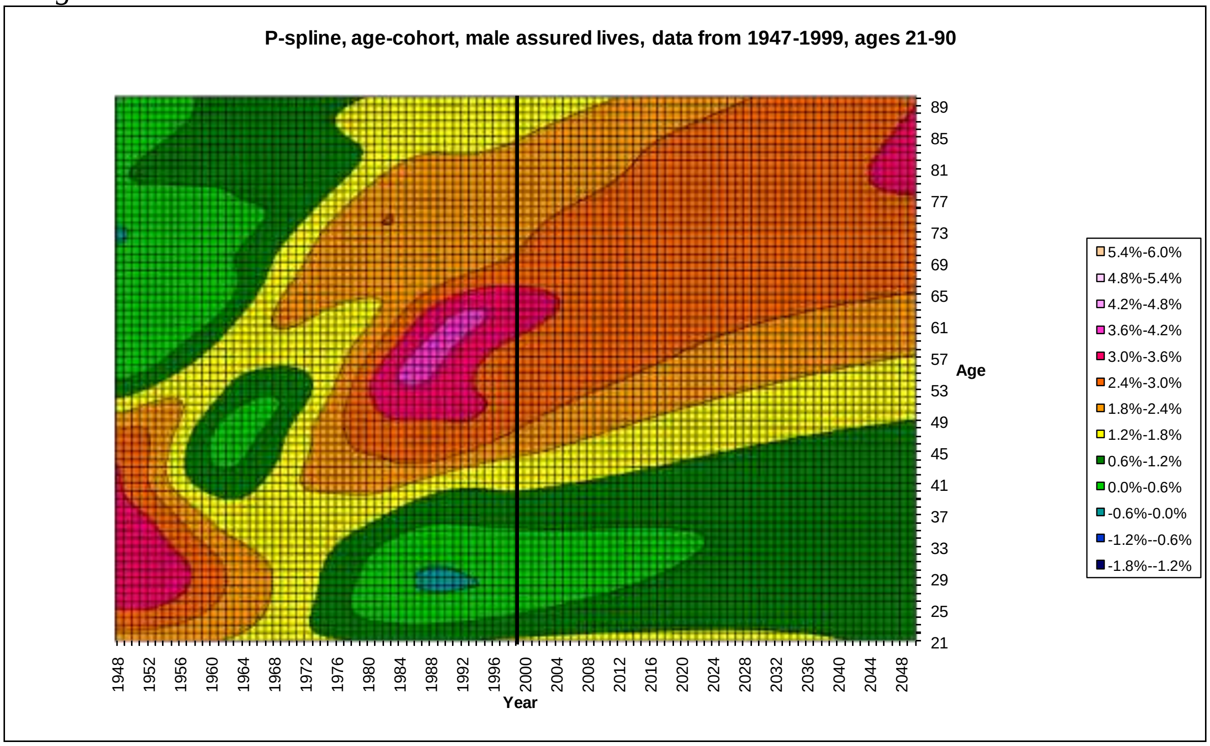

Figure B10

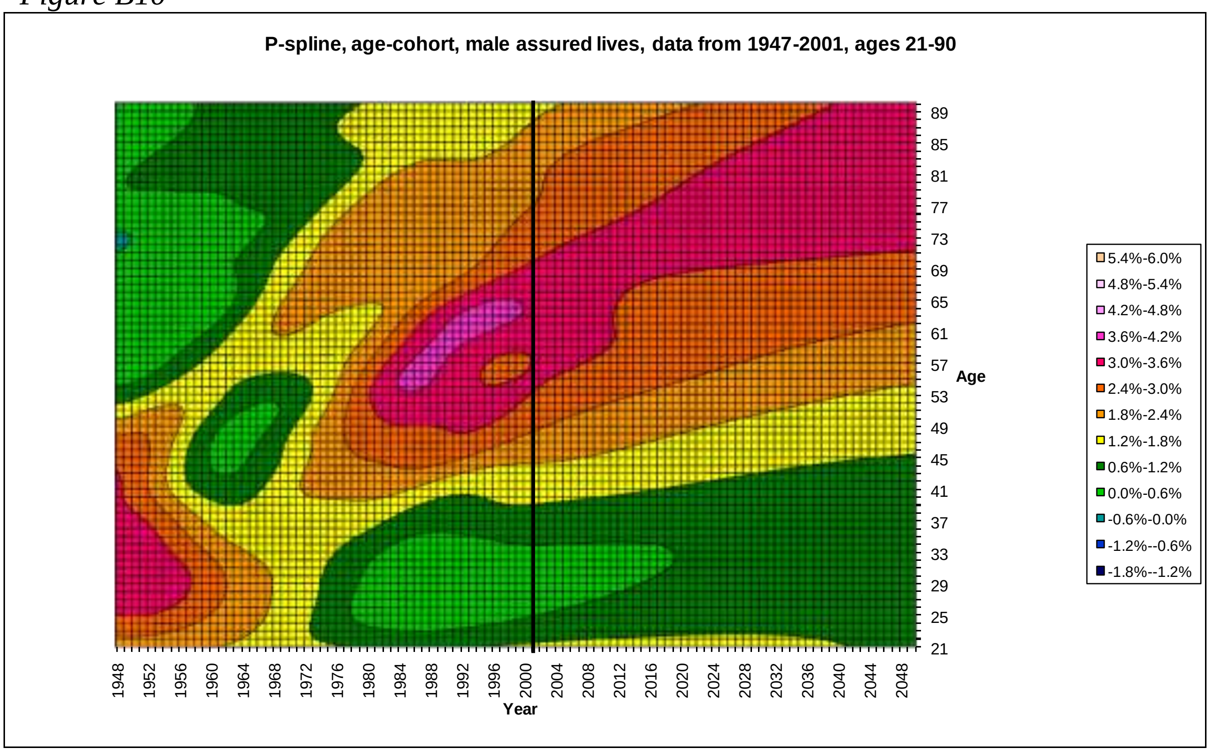

Figure B11

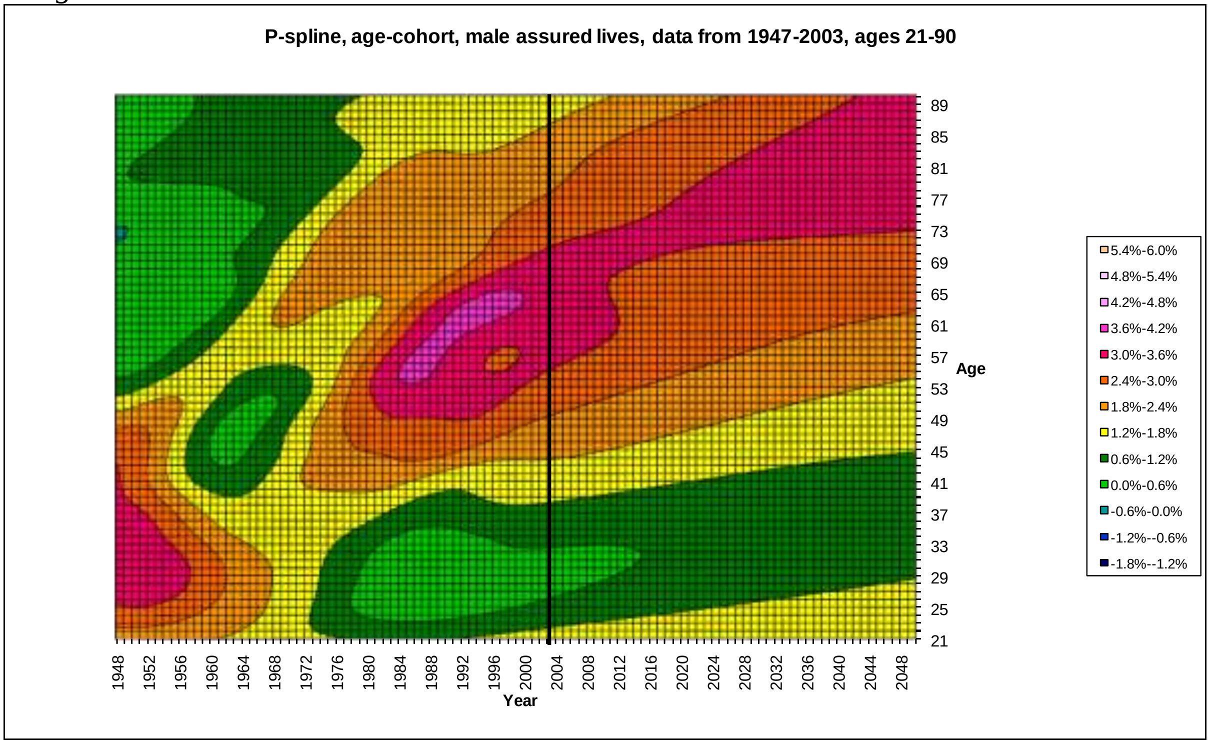
P-spline, age-cohort, male assured lives, data from 1947-2003, ages 21-90
Age
89
85
81
77
73
69
65
61
57
53
49
45
41
37
33
29
25
21
Year
1948
1952
1956
1960
1964
1968
1972
1976
1980
1984
1988
1992
1996
2000
2004
2008
2012
2016
2020
2024
2028
2032
2036
2040
2044
2048
5.4%-6.0%
4.8%-5.4%
4.2%-4.8%
3.6%-4.2%
3.0%-3.6%
2.4%-3.0%
1.8%-2.4%
1.2%-1.8%
0.6%-1.2%
0.0%-0.6%
-0.6%-0.0%
-1.2%--0.6%
-1.8%--1.2%

Appendix C – Contour maps of projected improvements for male assured lives using age-period penalties

Figure C1

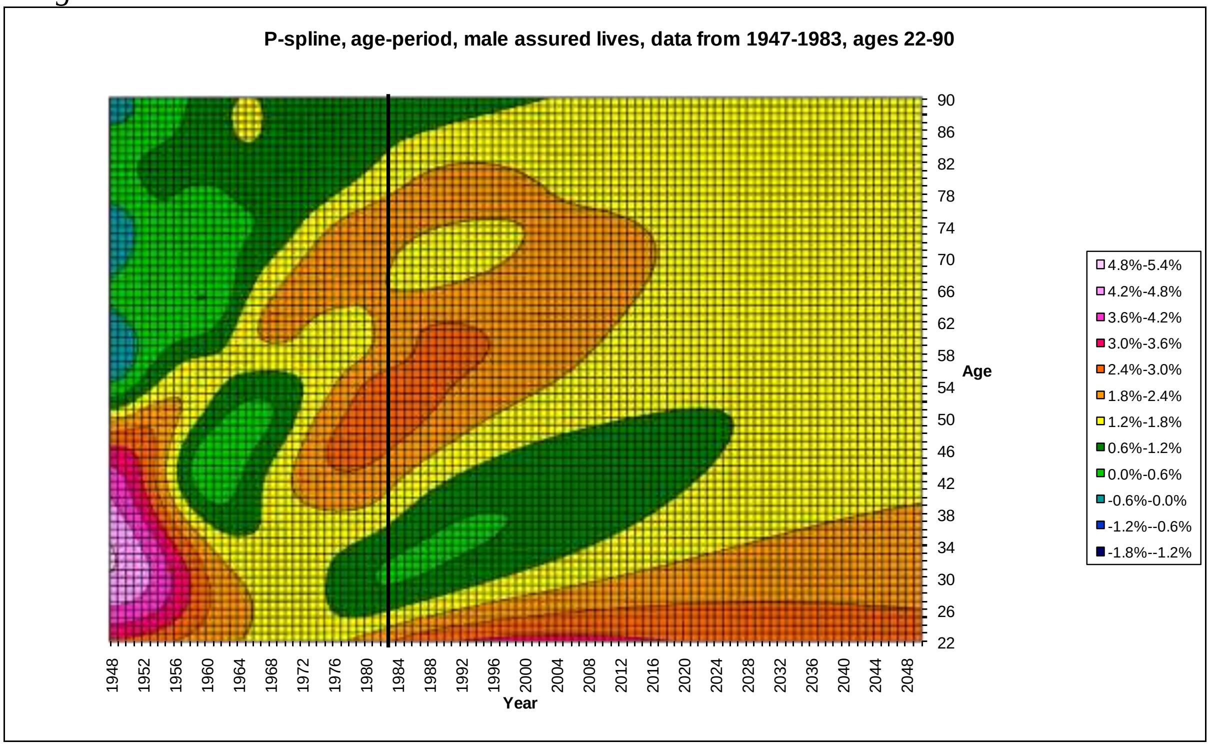

Figure C2

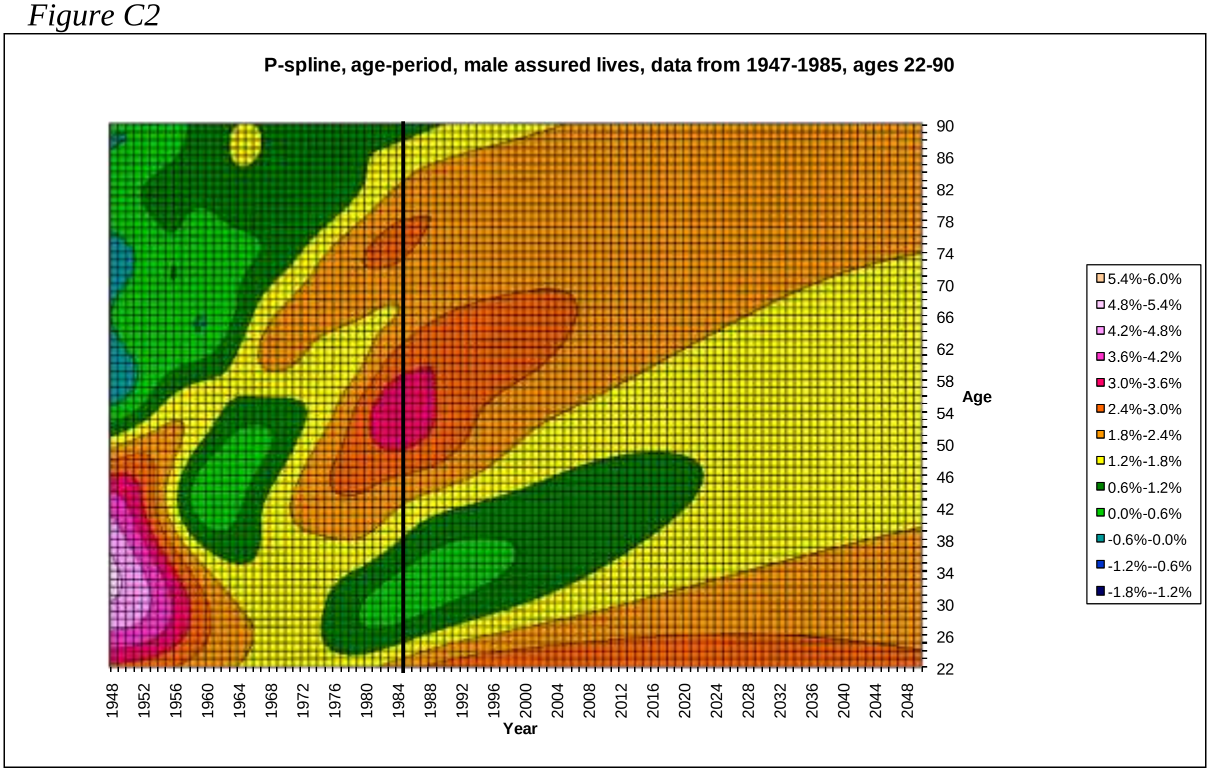

Figure C3

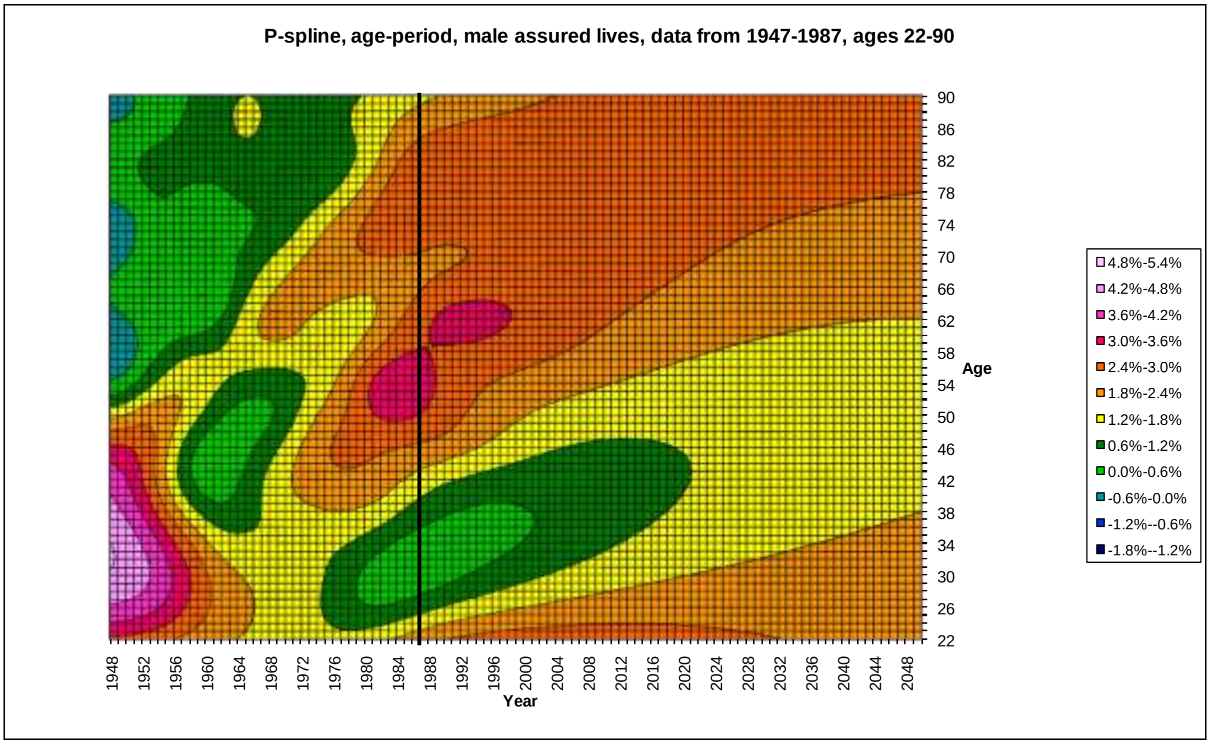

Figure C4

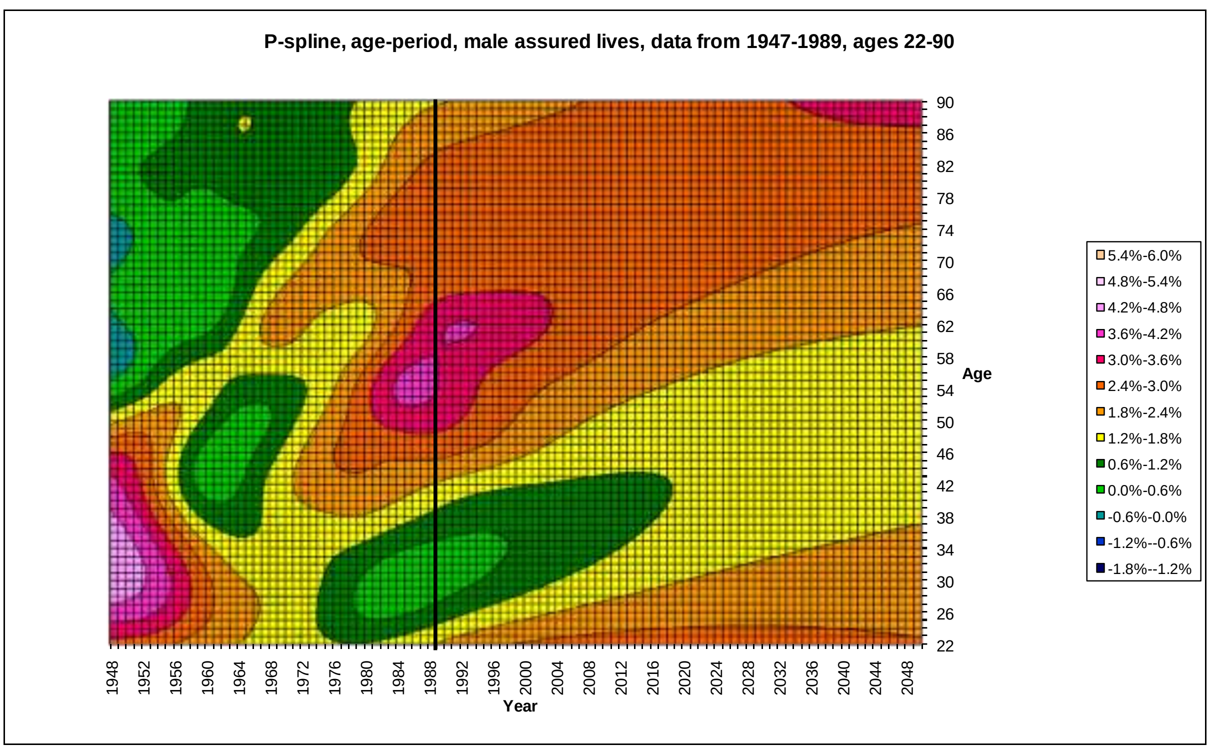

Figure C5

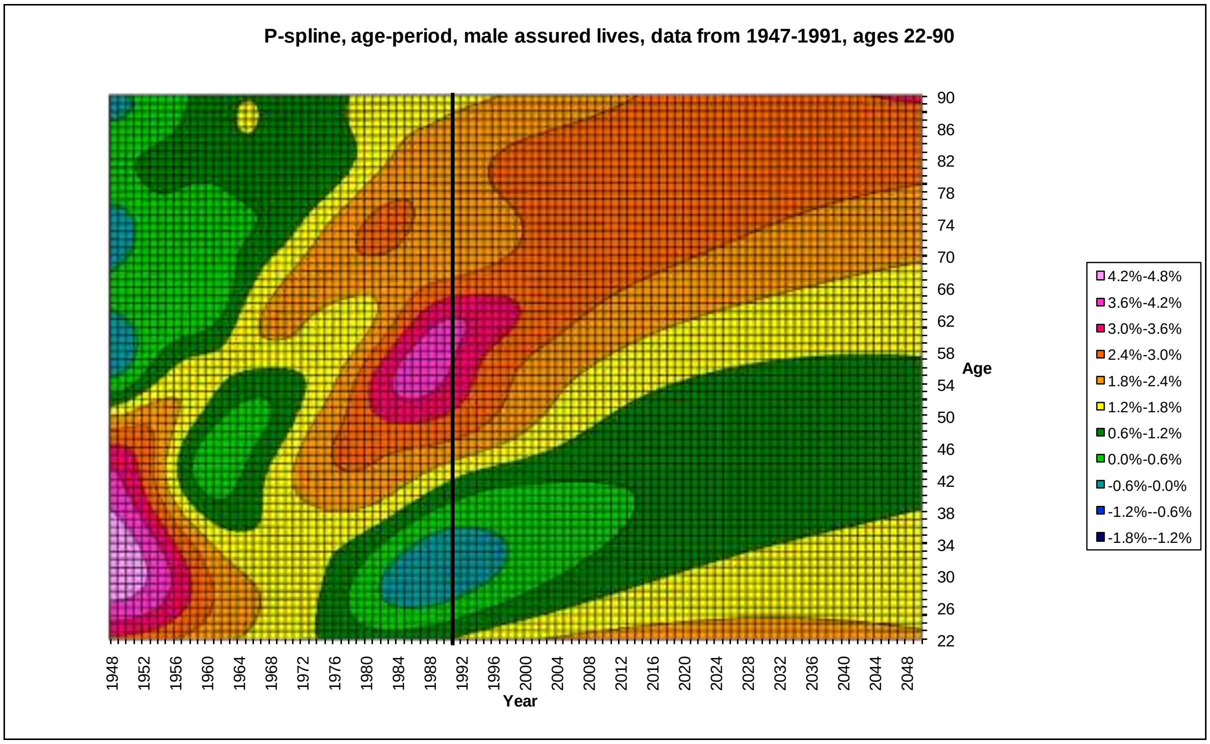

Figure C6

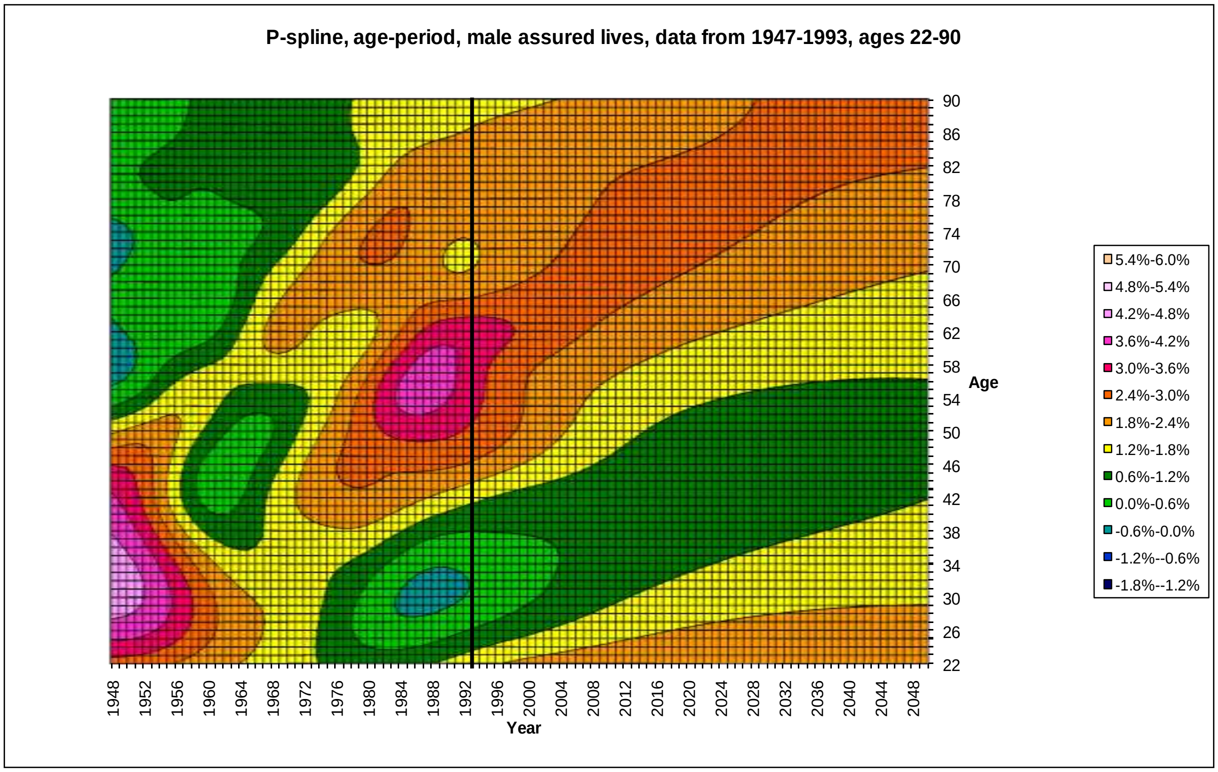

Figure C7

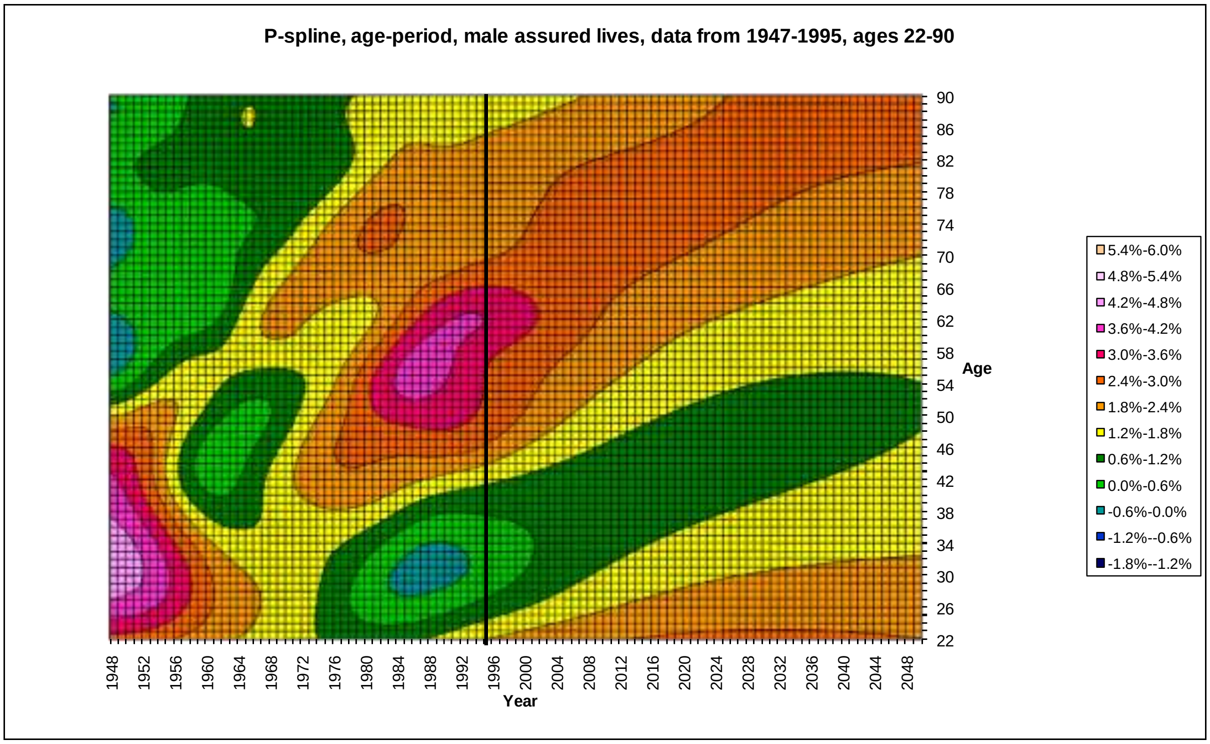

Figure C8

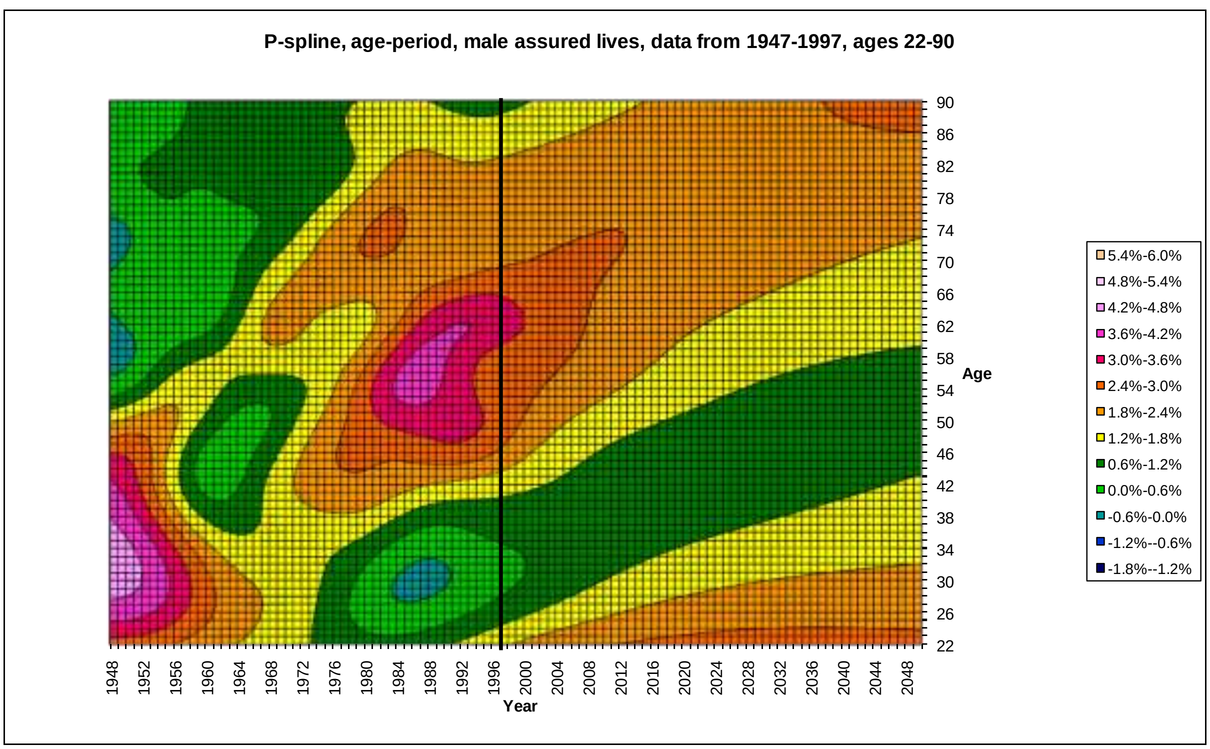

Figure C9

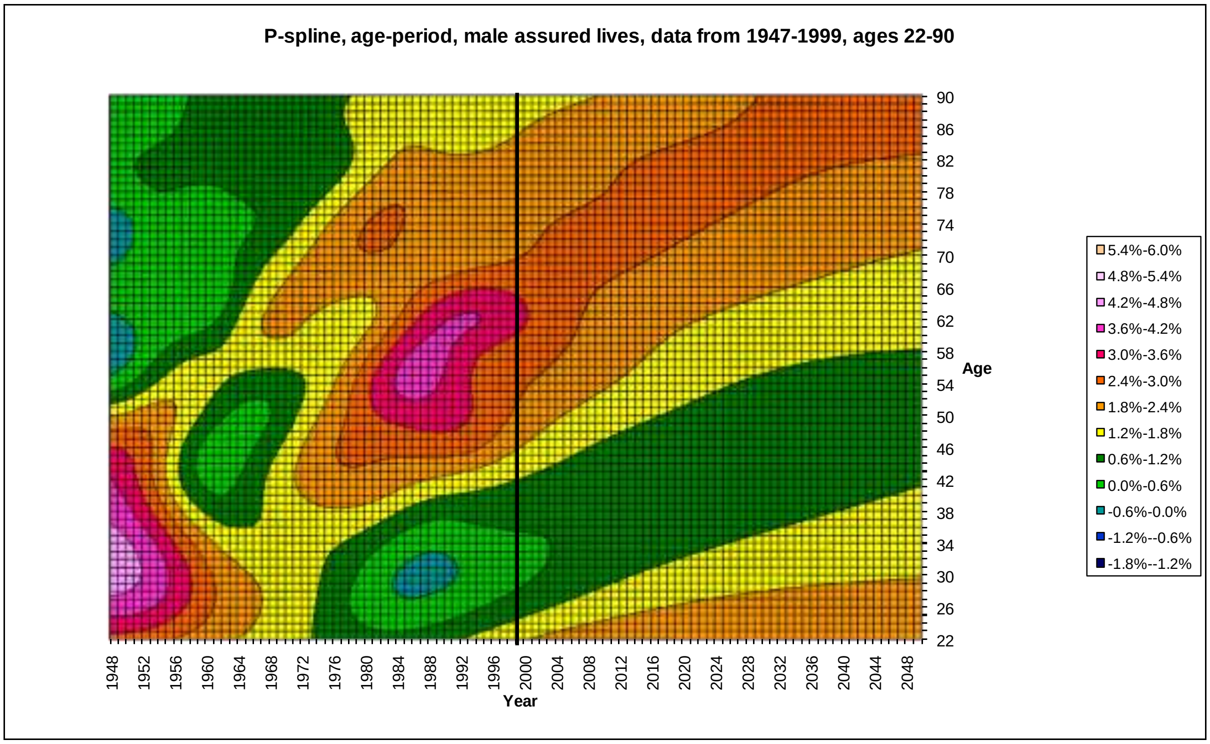

Figure C10

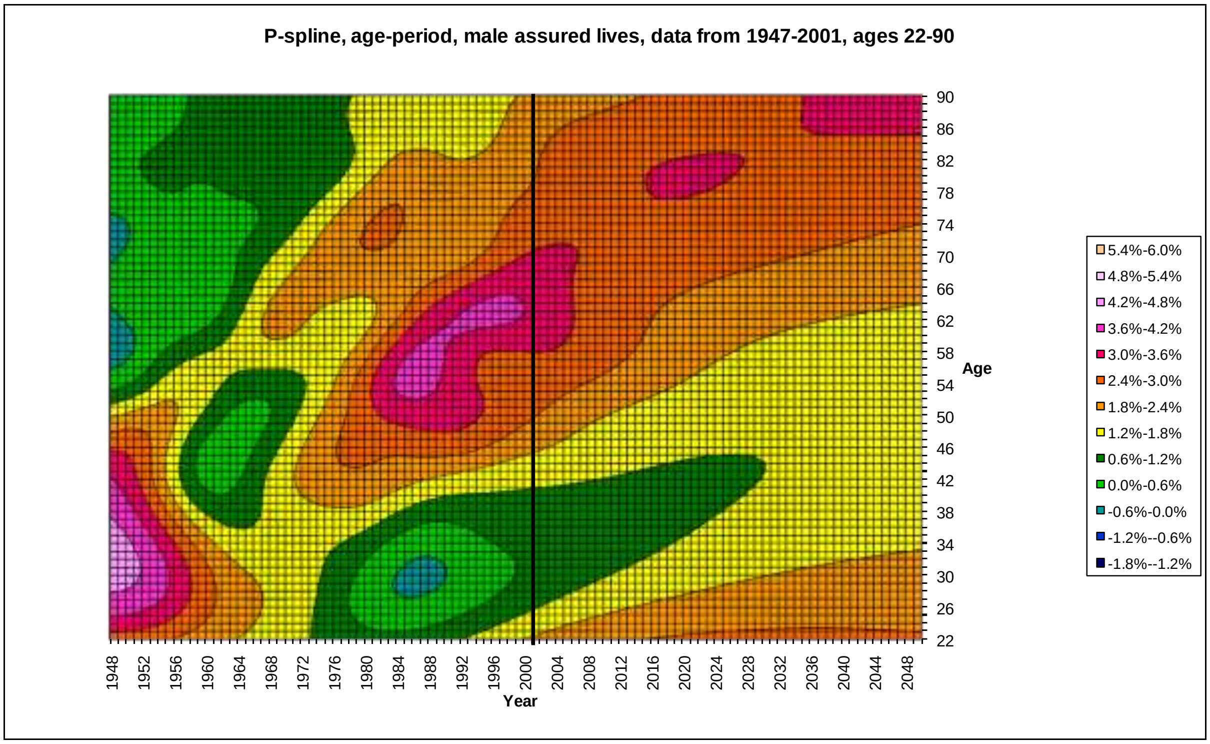

Figure C11

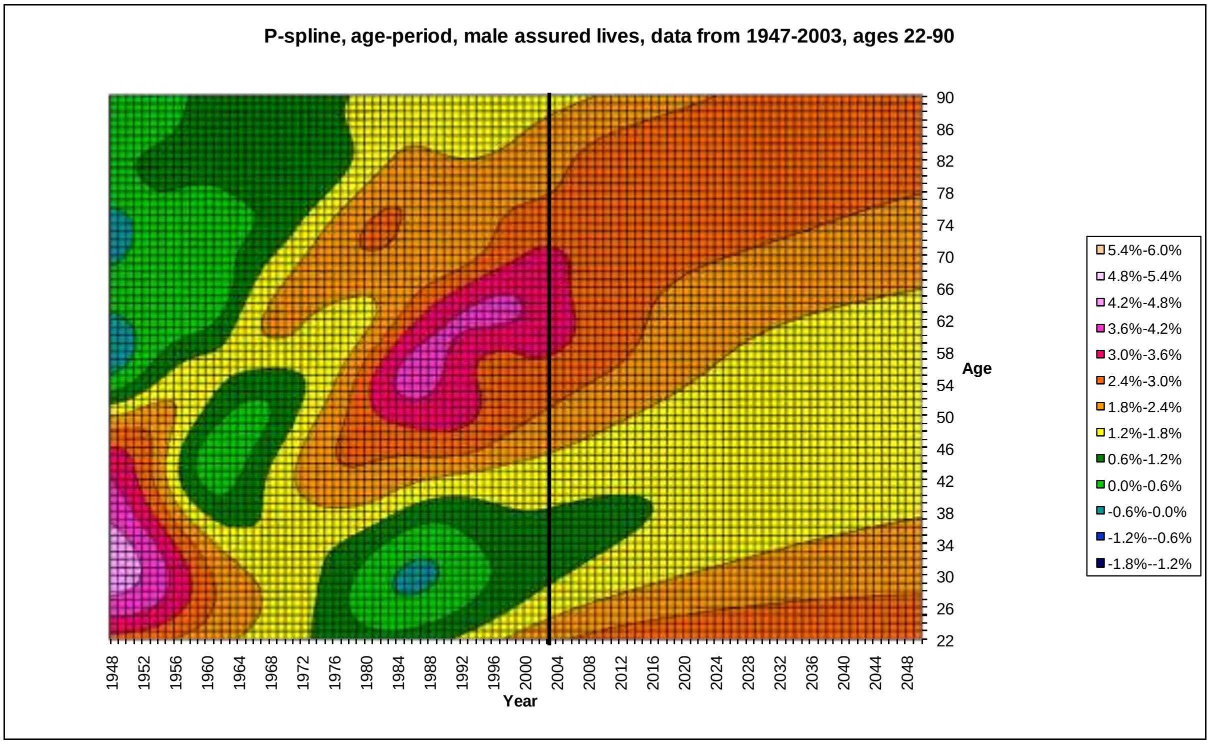
P-spline, age-period, male assured lives, data from 1947-2003, ages 22-90
90
86
82
78
74
70
66
62
58
54
50
46
42
38
34
30
26
22
Age
1948
1952
1956
1960
1964
1968
1972
1976
1980
1984
1988
1992
1996
2000
2004
2008
2012
2016
2020
2024
2028
2032
2036
2040
2044
2048
Year
5.4%-6.0%
4.8%-5.4%
4.2%-4.8%
3.6%-4.2%
3.0%-3.6%
2.4%-3.0%
1.8%-2.4%
1.2%-1.8%
0.6%-1.2%
0.0%-0.6%
-0.6%-0.0%
-1.2%--0.6%
-1.8%--1.2%

Appendix D – Effect of reducing the age range on projected improvements: contour maps of projected improvements for male assured lives using age-cohort penalties

Figure D1 - Ages 21 to 90

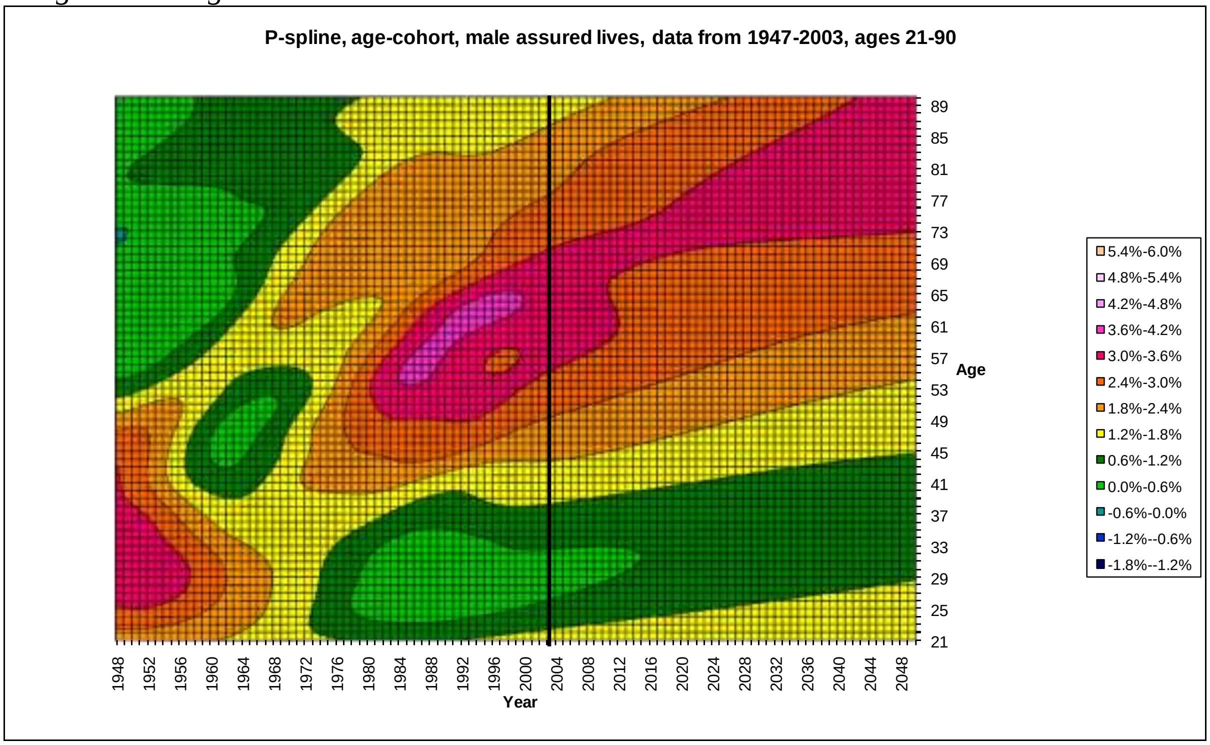

Figure D2 - Ages 42 to 90

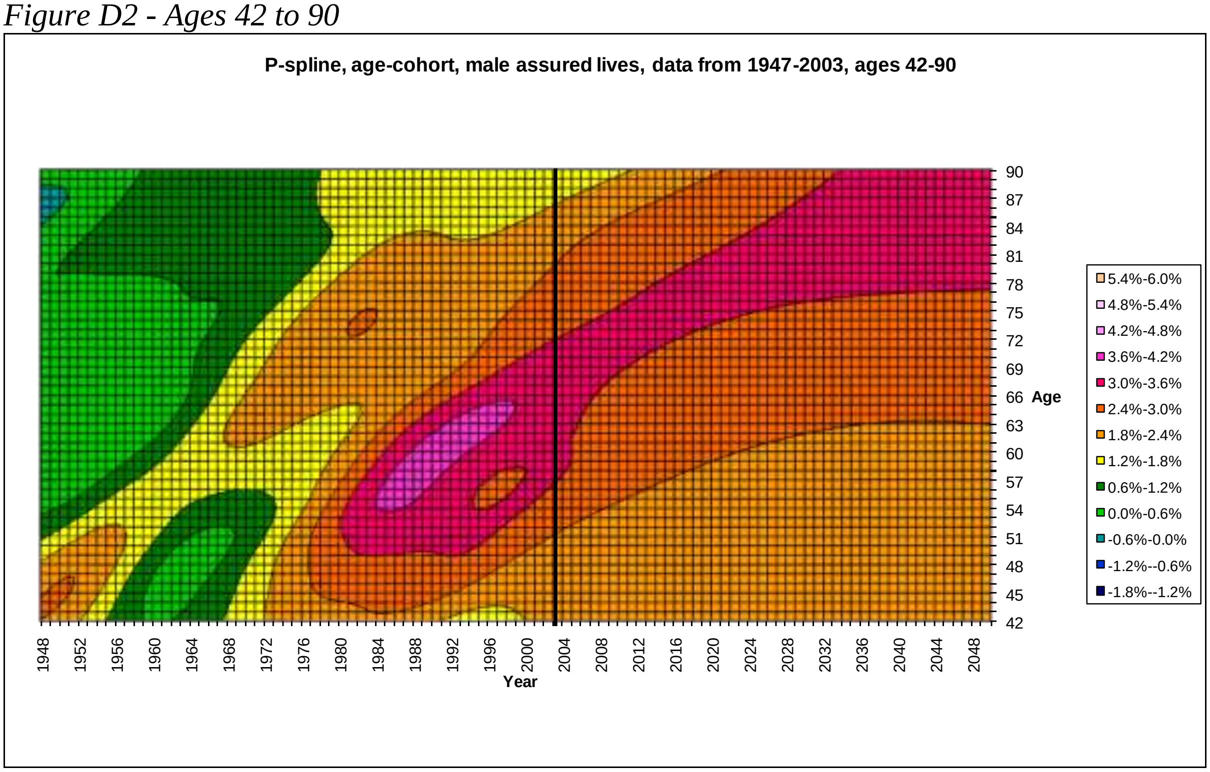

Appendix E – Contour maps of projected improvements for male ONS lives using age-cohort and age-period penalties

Figure E1 – Age-cohort penalty

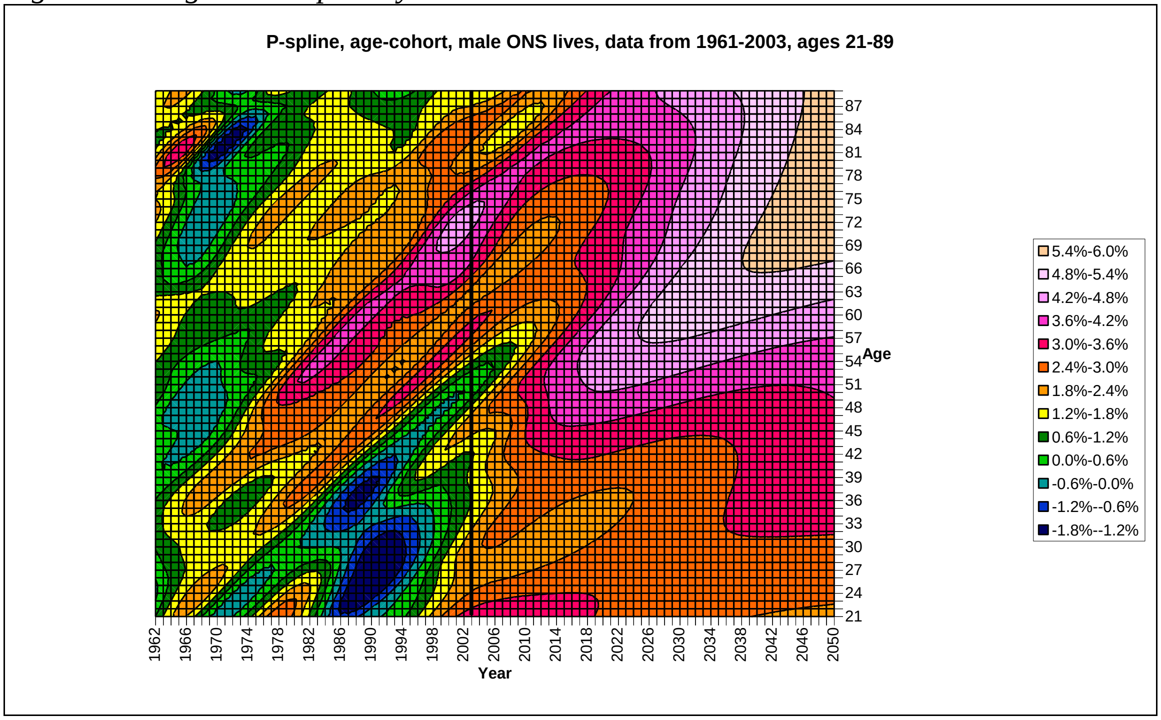

Figure E2 – Age-period penalty

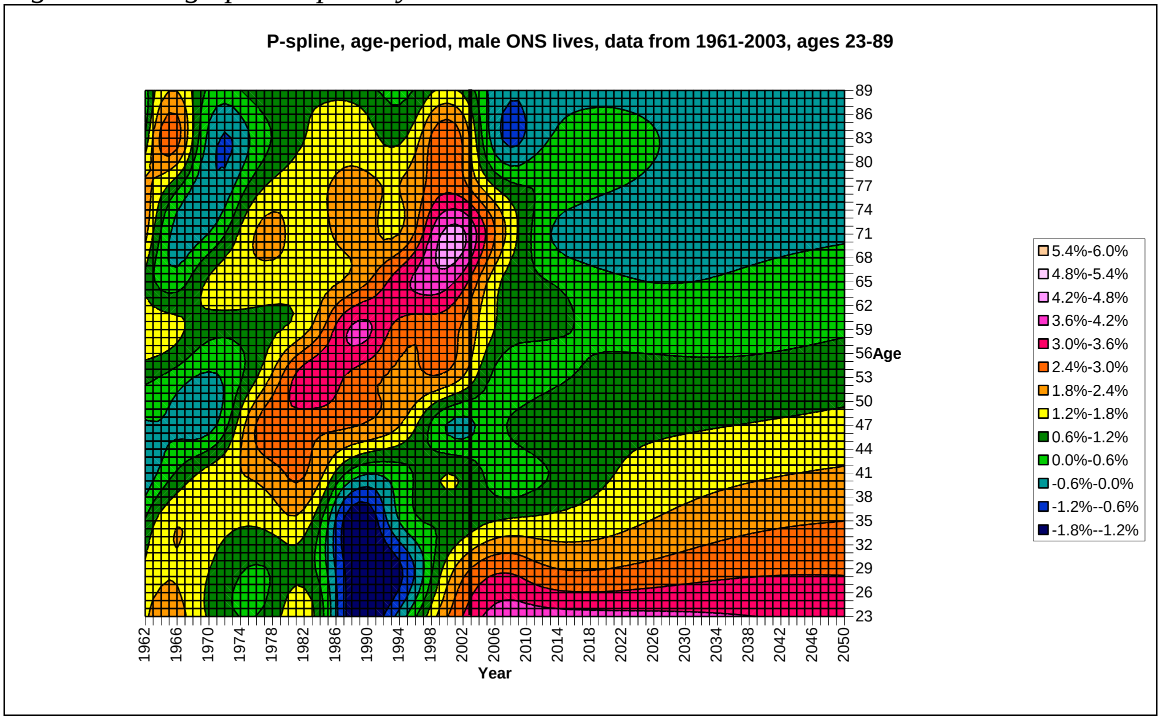

REFERENCES

CMI (2002). *Working Paper 1: An interim basis for adjusting the "92" Series mortality projections for cohort effects*. The Faculty of Actuaries and Institute of Actuaries.

CMI (2004). *Working Paper 3: Projecting future mortality: A discussion paper*. The Faculty of Actuaries and Institute of Actuaries.

CMI (2005a). *Working Paper 11: Responses to Working Paper 3 entitled Projecting future mortality: A discussion paper and further commentary thereon*. The Faculty of Actuaries and Institute of Actuaries.

CMI (2005b). *Working Paper 12: The graduation of the CMI 1999-2002 mortality experience: Feedback on Working Paper 8 and proposed assured lives graduations*. The Faculty of Actuaries and Institute of Actuaries.

CMI (2005c). *Working Paper 15: Projecting future mortality: Towards a proposal for a stochastic methodology*. The Faculty of Actuaries and Institute of Actuaries.

CMI (2005d). *Working Paper 16: The Graduation of the CMI 1999-2002 Mortality Experience: Proposed Annuitant and Pensioner Graduations*. The Faculty of Actuaries and Institute of Actuaries.

Olivier, P. & Jeffery, A. J. (2004) *Stochastic Mortality Models*. Presentation to the Society of Actuaries in Ireland 20/10/2004.

Renshaw, A & Haberman, S. (2005). *Mortality Reduction Factors Incorporating Cohort Effects.* Cass Business School Actuarial Research Paper No. 160

Richards, S.J. Kirkby, J.G & Currie, I.D. (2005). *The importance of year of birth in two-dimensional mortality data.* Presented to the Institute of Actuaries, 24 October 2005.

Willets, R.C. (1999). *Mortality in the next Millennium*. Staple Inn Actuarial Society, London.

Willets, R.C. (2004). *The cohort effect: Insights and explanations*. British Actuarial Journal **10**, 833-877.

Yang, S (2001). Reserving, pricing and hedging for guaranteed annuity options. PhD thesis, Heriot Watt University.