Continuous Mortality Investigation

Working Paper 39

A Prototype Mortality Projections Model:

Part Two – Detailed Analysis

It remains the responsibility of any actuary or other person using a projection of future mortality to ensure that it is appropriate for the particular purpose to which it is put, regardless of the source of the projection.

July 2009

Contents

1. Introduction

In June 2009 the Continuous Mortality Investigation (CMI) published a prototype mortality projections model with an accompanying paper ('Working Paper 38: A Prototype Mortality Projections Model: Part One – An Outline of the Proposed Approach') and a User Guide for the Model.

Working Paper 39 is designed to be read in conjunction with this earlier paper and does not repeat sections such as the background, the composition of the Working Party, the acknowledgements and the consultation process.

This paper provides more detail on the research carried out by the Working Party related to the structure of the Model and the choice of default parameter values for the 'Core' parameter layer. The structure of the remainder of this paper is as follows:

Section 2 outlines how the thinking of the Working Party evolved over time, linking through from the research to the proposed design of the Model.

Section 3 outlines issues relating to the choice of Initial Rates of Mortality Improvement for the 'Core' parameter layer of the Model.

Section 4 describes research that could inform the choice of Long-term Rates of Mortality Improvement.

Section 5 provides detail on the convergence methodology and choice of 'Core' layer parameter values relating to convergence from Initial to Long-Term Rates of Mortality Improvement.

Section 6 illustrates sample projections and the sensitivity of the output of the Model to varying selected parameter values. A spreadsheet containing the results of additional sensitivity tests is available from the CMI pages on the UK Actuarial Profession's website.

2. How the Thinking of the Working Party Evolved Over Time

The Working Party felt that understanding of the research described in this Working Paper might be aided if the development of its thinking over time was briefly described.

Given its primary objective was to produce a model which was relatively straightforward to understand, use and modify (with similar perceived advantages to the Interim Cohort Projections), the Working Party quickly discounted more complex options (such as cause-based, disease-based or stochastic methodologies). The structure of assuming 'current' rates of change blending into 'long-term assumptions' – adopted by the GAD/ONS for recent National Population Projections – seemed sensible given these aims.

The first task was therefore to identify a methodology suitable for deriving Initial Rates of Improvement. Given the work of the CMI Mortality Projections Working Party in recent years – summarised in section 1.1 of Working Paper 38 – an obvious first step was to explore using a P-Spline model to smooth historic experience.

However, the Working Party was concerned that adopting a P-Spline methodology could introduce unwanted 'edge effects', i.e. that the Initial Rates of Improvement for the final year of the data would be unduly influenced by the relative level of experience in that final year and hence would vary materially from one year to the next. Section 3.1.2 describes how this potential volatility appears to be controlled by 'stepping back' two years from the last year of the dataset. For instance, for the prototype model, Initial Rates of Mortality Improvement are estimated for calendar year 2005 using a dataset to calendar year 2007.

The P-Spline approach, with this modification, was then compared with alternative methods of deriving Initial Rates of Improvement. This work, described in section 3.1.3, concluded that age-cohort P-Spline models fitted to England & Wales population data produced reasonable estimates of 'current' improvement rates in that dataset.

The Working Party then considered what data could be used to parameterise the Model. Initially, the idea of incorporating suitable parameters in the Model to allow suitable transformations to default rates of improvement derived using population-level data was explored. This idea stemmed from the observation that historic rates of improvement have been somewhat more rapid for higher socio-economic class groups than the population average (Willets *et al*, 2004). It also reflected the research, described in CMI Working Paper 1, which showed an earlier birth cohort experiencing peak improvements in CMI data versus general population data. Specifically, the CMI data then showed those born in 1926 experiencing the most rapid improvements, whereas population data showed the 1931 birth cohort demonstrated the most rapid rates of change. Furthermore, researchers in other fields have described the 'diffusion theory of behaviour change', where changes in health risk behaviour are adopted first among the middle classes and then diffuse through the population (e.g. Evandrou & Falkingham, 2002).

Section 3 of this paper compares improvements observed in the general population with those seen in alternative datasets, i.e. CMI datasets for assured lives and life office pensioners; data supplied by Club Vita for pensioners in self-administered pension schemes and ONS data for different socio-economic class groups.

The Working Party's analysis confirmed previous findings that England & Wales population data showed clear evidence of persistent year-of-birth cohort features – with the strongest feature peaking at the 1931 cohort – together with a more general increase in mortality improvement rates across a wide range of ages over the last 25 years. However, patterns of improvement by year-of-birth appeared somewhat different for the other datasets and the evidence to support different peak year-of-birth cohorts by socio-economic group or for insured/pensioner/population groups did not show a clear-cut, easy to interpret, pattern.

It was therefore concluded that the Model should be parameterised with reference to England & Wales population data, whilst leaving users with the flexibility to incorporate their own rates of initial improvement should they see fit.

A further point of model structure, which required consideration by the Working Party, was whether the Initial Rates of Mortality Improvement should be projected into the future 'by age' or 'by birth cohort.' The approach adopted by GAD/ONS using a similar method has been to project improvements 'by birth cohort' for some birth years and 'by age' for others.

One drawback of this approach is that it leaves missing 'triangles' in a matrix of future years and ages that need to be populated using an alternative method. The Working Party felt a more robust methodology would be to split improvements into two components, those due to 'age and period' and those due to 'year of birth.' An advantage with this approach is that rates of improvement due to the different components can converge towards long-term assumptions over different periods and at different speeds.

Section 3.4.2 describes how an age-period-cohort model was constructed to split the improvements into Age/Period and Cohort Components.

Once Initial Rates of Mortality Improvement had been derived, the Working Party considered what approach to adopt for the convergence towards the long-term assumptions. The proposed methodology, which utilises a series of cubic polynomials is described in section 5. As discussed in Working Paper 38, the methodology was selected so that it would have the flexibility to generate scenarios in which the rate of change continues to accelerate in the short-term, before decelerating in the longer term. The Working Party's analysis of past features of mortality improvements led it to introduce flexibility in the Model to vary periods of convergence by age and by cohort.

The final element of the model structure is the long-term rate of mortality improvement. The Working Party concluded that the choice of the Long-Term Rate of Mortality Improvement is necessarily subjective and is the single most important parameter for Users to set for each projection. No default value is proposed in the prototype Model, however some analysis that may help Users is summarised in section 4.

The sensitivity of projections to changes in parameter values was investigated and used to guide the Working Party in designing the separation of 'Core' and 'Advanced' parameter layers. The Core level parameters selected by the Working Party are those to which the results are most sensitive and for which Users may more readily form a view.

The Working Party had then to decide what default values to assign to parameters required for Core Projections. The rationale for the choice of default values for the Model is outlined where appropriate during the course of this paper.

3. Initial Rates of Mortality Improvement

3.1. Estimation of Current Rates of Mortality Improvement

3.1.1. Data

The Working Party sought estimates and insights on past and present mortality improvements through analysis of data both at population level and at the level of large pools of insured lives and pensioners. Analysis was concentrated on the three datasets summarised in the following table:

	Population		**Assured Lives**		**Pensioners**	
Source	ONS		CMI		CMI	
Coverage	England & Wales Population		Permanent Assurances (with contributing UK Insurers) All durations		Pensioners (with contributing UK Insurers) All durations	
Age Range - dataset - used	0 to 104 18 to 102		20 to 90 20 to 90		20 to 100 60 to 100	
Time Period	1961 to 2007		1947 to 2006 (M) 1983 to 2006 (F)		1983 to 2006	
Death Counts	Deaths Registered		Settled Death Claims, notified within 6 months of year-end		Deaths Notified within 6 months of year-end	
Average per Year [1993-2007; Thousands]	Males	Females	Males	Females	Males	Females
Total for Ages 30-49	13.32	8.24	1.05	0.40	0.00	0.00
Total for Ages 50-69	62.21	40.10	5.56	1.37	0.90	0.28
Total for Ages 70-89	156.82	172.95	4.14	1.26	9.54	2.20
Exposure Measure	Mid-year population estimates (ONS)		Life years exposed-to-risk (census method)		Life years exposed-to-risk (census method)	
Average per Year [1993-2007; Millions]	Males	Females	Males	Females	Males	Females
Total for Ages 30-49	7.464	7.554	0.986	0.547	0.000	0.000
Total for Ages 50-69	5.500	5.714	1.069	0.388	0.062	0.038
Total for Ages 70-89	2.308	3.375	0.096	0.046	0.147	0.052

Table 3.1: Principal datasets used for analysis of past rates of mortality improvement

All three datasets provide counts of deaths, and some measure of exposure, for 1-year cells by individual age and calendar year, and provide separate data for males and females. This level of detail in the data is sufficient to support the application of statistical tests and models, such as P-Splines, in the analysis of the data.

The simple data volume measures shown in Table 3.1 give a broad indication of the relative size of the datasets and of their distribution by age and gender. It should also be noted that mortality trends observed from the CMI (and other pooled) datasets may be distorted by changes over time to the mix of contributing offices (or pension schemes).

In addition the Working Party is grateful to Club Vita for supplying results from their analysis of self-administered pension scheme data submitted to the Club Vita database. This gave the Working Party access to both crude and smoothed (via age-cohort P-Spline) annual rates of mortality improvement, separately for males and females, for 1-year cells across the age range 50-100 and calendar years 1994-2007 along with summary statistics of total exposure for each calendar year, and total exposure for each age.

The scale of the dataset underlying the Club Vita results supplied is illustrated by exposure measures consistent with those shown in Table 3.1. The average yearly exposure for male lives over the period 1993 to 2007 was 0.163 million life-years for ages 50-69, and 0.128 million life-years for ages 70-89; the corresponding figures for female lives are 0.164 million and 0.160 million.

3.1.2. Avoiding Undue Influence of 'Edge Effects'

Given stochastic and other volatility inherent in the data, it is necessary to smooth the observed mortality experience, across both age and time, in order to see the underlying patterns of mortality improvements more clearly.

The Working Party initially favoured using P-Spline methodology (which is documented, for example, in CMI Working Papers 15 and 20) to smooth historic experience. P-Spline models were applied to the deaths and exposure data to produce fitted, smooth, two-dimensional mortality rate surfaces, across age and time, in the region of the data for each gender; the projection elements of P-Spline models were not used in this work.

However, before progressing the analysis, the Working Party addressed two important concerns:

1. That P-Spline methodology (or indeed other smoothing methodologies) may be vulnerable to 'edge effects': that is, that the estimated rates of mortality improvement for the final year of the data would be unduly influenced by the relative level of experience in that final year, and hence could vary materially from one year to the next.
2. That the structure of P-Spline models may emphasize the strength of certain patterns in the data beyond the level statistically justified by the raw data.

These two issues are considered in this and the next section respectively.

The concern over potential edge effects was investigated by back-testing a series of P-Spline models applied to ONS data for the population of England & Wales. P-Spline models were fitted in pairs to data for overlapping time-periods, and the resulting surfaces compared to evaluate roughly the risk of distortion in estimating improvements at the leading edge (i.e. the last year) of available data.

For example, a pair of models was fitted to data for the periods 1961-1970 and 1961-1979, and then the derived rates of mortality improvement for 1970 were compared to see how the first estimate, where 1970 marked the edge of the data, differed from the second estimate, where 1970 lies in the middle of the data so that the fitted mortality surface takes account of post-1970 experience.

Pairs of P-Spline models were fitted in this way to a series of time intervals, varying in length from 8/16 years to 20/40 years and covering the whole period available through the England & Wales population dataset.

This exercise clearly demonstrated that edge effects can lead to material differences in estimates of mortality improvement rates for single calendar years. There was a significant number of pairs for which the comparison of estimates, from the leading-edge and the middle year, showed absolute differences in mortality improvement rates of more than 1% p.a., although other pairs resulted in a much closer match between the two estimates.

Study of the residual errors from the fitted models (i.e. the differences between the fitted and observed mortality rates) gives a key insight into this problem. A heat map of the residuals, measured in terms of standard errors, reveals significant calendar-year peaks and troughs, applying across a wide range of ages, rather than a random distribution of residual errors which would have been expected if the only volatility were due to sampling from a stochastic process. The overall, underlying level of mortality varies significantly from year to year but these peaks and troughs are smoothed out by the P-Spline model (and indeed by other methodologies which smooth over time).

The pairs of models which gave concern – for years where the leading-edge estimate differed materially from the middle-year estimate – generally related to years which (with hindsight) may be recognised as having been relatively heavy or light in terms of mortality experience. P-Spline models generally produce surfaces which fit the data more closely at the edges than across the middle, as there is relatively greater parameter freedom at the edges. So if, for example, a year with relatively heavy mortality lies on the leading edge of the data, the P-Spline model adheres closely to the overall level of mortality for that year and the estimated mortality improvement rates for that year will be relatively low. However, when further years' data is added, the model is better able to 'recognise' experience for the year as relatively heavy, and the fitted surface smoothes across years, and so will lead to a generally higher set of estimated mortality improvement rates for that year.

The Working Party's analysis showed that these edge effect problems arise similarly for both age-period and age-cohort P-Spline models, and that they are not satisfactorily resolved by moving the knot position within the P-Spline model or by varying the spacing of knots.

As the problem is one of difficulty in judging whether the latest year's experience will ultimately be seen to have been above or below the trend line, the Working Party considered whether it might be possible to use other data – such as cause of death statistics – to improve that judgement, so that mortality improvement rate estimates could then be adjusted in some way. However, it was concluded that this would not be a practical way forward.

Instead, the Working Party investigated the idea of taking the estimate of mortality improvement rates from a time-point set back from the leading edge. Further analysis clearly showed that the differences between the two sets of estimates were significantly reduced by taking the estimate from inside the edge of the data. Moving further inside by one year at a time reduced the differences each time, although such gains in reliability diminished with each successive step.

Mortality improvement rates were estimated for each available year, 1971-2006, by applying a P-Spline model to the previous 10-years' data (i.e. 1961-71 for 1971), and these leading-edge estimates were compared to estimates derived using the entire ONS dataset period (1961-2007); the mean (averaged across all ages and years) absolute difference in rates was 0.50% p.a. This measure of estimation error reduced to 0.40% if the first estimate was taken 1 year inside the edge of the data, and to 0.32%, 0.26% and 0.22% for estimates taken 2, 3 and 4 years inside the edge respectively. Similar results were obtained for estimates derived from 20-year datasets.

The Working Party concluded that mortality improvement rates derived on the leading edge of the data are subject to too great an estimation error, and that the issue could be best managed by taking estimates from inside the edge. A trade-off is required between estimation uncertainty and currency. The Working Party opted to base its estimates on improvement rates calculated 2 years inside the edge of the data, effectively using the last two years' data to allow a small amount of hindsight into the calculation so that the estimates are less likely to be significantly revised when further years' data are added as the experience unfolds.

3.1.3. Comparison of Estimators for Current Rates of Mortality Improvement

The Working Party next addressed the concern that the structure of P-Spline models may emphasize the strength of certain patterns in the data beyond the level statistically justified by the raw data. The Working Party considers that there is no uniquely correct way of estimating mortality improvement rates and that each approach has its own merits and potential for error.

As the problem is akin to model error, it was decided that the best way forward was to present mortality improvement rate estimates derived from a number of alternative methodologies. By comparing the different estimates, views may be formed on:

- The range of estimated values, as a crude measure of uncertainty; and
- The strength of particular features in the results, distinguishing between those (strong) features which emerge regardless of methodology, and those (questionable) features which appear to be methodology-dependent.

The approach is illustrated in this section using the ONS dataset for the male population of England & Wales. Mortality improvement rates were estimated for 2005, i.e. 2 years inside the edge of the available data.

As a first step, alternative P-Spline models may be compared: age-period and age-cohort models were fitted to the data, and the spacing of knots was varied between 3 and 6 years. Where the knot-spacing is smaller the model has more parameters and so is likely to fit more closely to the data, subject to the trade-off between closeness-of-fit and smoothness of the fitted surface (as governed by the smoothing penalty function).

Figure 3.1 shows the rates of mortality improvement for 2005 derived from age-period (AP) and age-cohort (AC) P-Spline models with 3, 4 and 5-year knot-spacing.

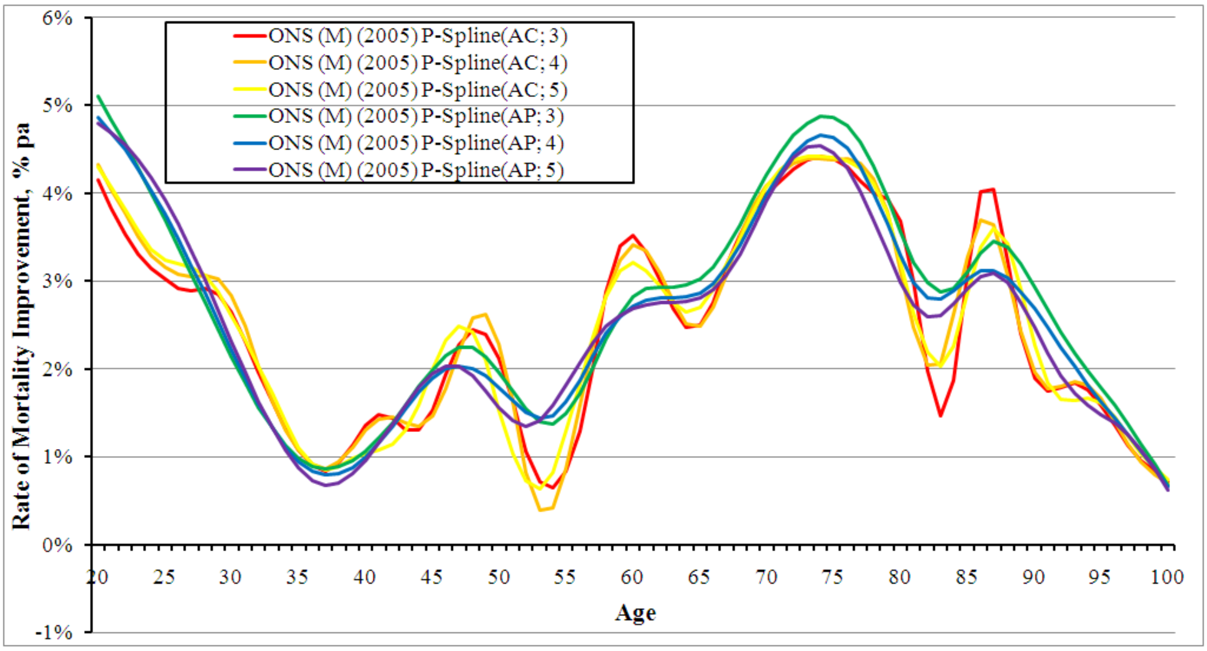

Figure 3.1: Annualised rates of mortality improvement in 2005: Comparison of estimates derived from alternative P-Spline models fitted to the ONS dataset for the male population of England &Wales

For this dataset, the largest of those used in terms of number of deaths, the following points may be noted:

- Overall variation between alternative P-Spline models is second order compared to the variations by age in rates of mortality improvement.
- Peaks and troughs in rates of mortality improvement appear to be slightly intensified by age-cohort models compared to age-period models.
- Closer knot spacing tends to produce a slightly more complex pattern, suggesting that 5-year spacing may over-smooth for the male population data set.

A number of alternative, crude methodologies were also applied to estimate rates of mortality improvement for 2005. It seems plausible to argue that the underlying mortality rates should be smooth by age within a calendar year, but that they may not be smooth between years (although, the observed rates will, of course, display volatility in both cases). Therefore, each of these methodologies involved some degree of smoothing of the observed mortality rates across age, independently for each calendar year. Annual rates of mortality improvement were calculated for each individual age and year, and then averaged over the 5-year period 2003-2007, either for a given age or year-of-birth cohort.

Three of the methods tried for smoothing individual calendar year data across age were:

- Fitting 1-dimensional P-Spline models to each year's data
  - Using 5-year knot-spacing
  - The penalty-function gives a trade-off between fit and smoothness
  - The fitting takes account of the data volume at each age, so that smoothness dominates over fit where the observed rates have less statistical credibility.

- Fitting a series of cubic basis-splines to each year's data
  - Using 5-year knot-spacing
  - Fitting is by a least-squares method applied to $\log(m_x)$ on observed rates; data volumes are not taken into account
  - There is no smoothing penalty, so nothing to shape the rates even where data volumes are low; smoothness comes only from the grid of smooth basis-splines.
- Calculating a crude running average of mortality rates
  - For example, calculating the rate for age x by averaging rates for ages x-4 to x+4.

Figure 3.2 compares the rates of mortality improvement for 2005 derived from age-period and age-cohort P-Spline models with 4-year knot-spacing against rates derived using these alternative methodologies.

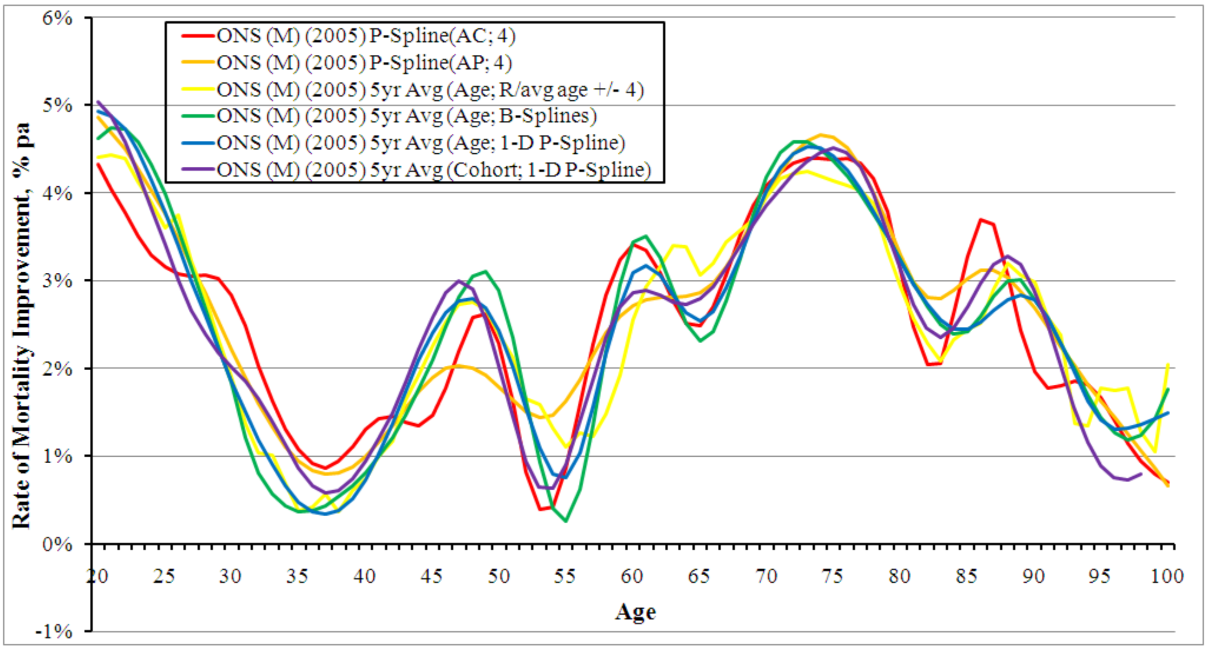

Figure 3.2: Annualised rates of mortality improvement in 2005: Comparison of estimates derived from fitted P-Spline models and various alternative methodologies applied to the ONS dataset for the male population of England &Wales

The following points may be noted, again based on the largest of the male datasets used:

- The correlation between P-Spline estimates and the alternative measures is good.
- Overall variation between estimates is second order compared to the variations by age in rates of mortality improvement.

Figure 3.2 also provides a convenient way to illustrate and assess estimate uncertainty. For this large dataset, several strong features of the data emerge consistently for all the measures. Uncertainty in the estimate of mortality improvements rates for this large dataset is generally of the order of ±0.5% p.a., and is smallest for ages 60 to 80 (where there are most deaths) and increases towards the youngest and oldest ages.

These conclusions on correlation of estimates and degree of estimation uncertainty are highly dependent on the dataset used (primarily due to differing data volumes) as will be seen in the next section.

3.2. Current Rates of Mortality Improvement

3.2.1. Population of England & Wales

The ONS dataset for the population of England & Wales contains data up to 2007. The Working Party considers that 2005, i.e. 2 years inside the edge of the available data, is the latest year for which sufficiently robust estimates of rates of mortality improvement may be made.

Estimates for male lives have already been presented in Figures 3.1 and 3.2. Estimates of rates of mortality improvement in 2005 for female lives, derived using the same set of methodologies, are shown in Figures 3.3 and 3.4.

The comparison of estimates of rates of mortality improvement for female lives confirms many of the observations drawn from the analysis of data for the male population:

- Several strong features of the data – peaks or troughs in the rates of mortality improvement by age – emerge consistently for all the measures when they are applied to the population datasets.
  - In particular, note the highest and broadest peak, centred on age 74 (year-of-birth 1931) for both males and females, …
  - … and other local peaks around ages 47 (1958), 61 (1944) and 87 (1918).
- The overall variation between estimates is generally second order, compared to the variations by age in rates of mortality improvement, for the population dataset.
- Whilst the overall variation between alternative P-Spline models is second order, peaks and troughs in rates of mortality improvement do appear to be slightly intensified by age-cohort models compared to age-period models.
- Closer knot spacing in the P-Spline models tends to produce slightly more complex patterns, suggesting that 5-year spacing may over-smooth for the population datasets.
  - However, the Age-Cohort P-Spline model with 3-year knot-spacing diverges markedly from the other estimates for the female population: this model is a poor fit to the data (and illustrates problems which may occur if the knots are too close).
- The correlation between P-Spline estimates and the alternative measures is generally good for the population datasets.

In addition, comparison of the charts for males and for females shows the variation in estimates, both between P-Spline models and against the alternative methodologies, is greater for the female population dataset than for the male population dataset.

Although clear features of the data for females do emerge consistently for all the measures for ages 55 to 80, there is less clarity at other ages. Uncertainty in the estimate of rates of mortality improvement for the female population dataset ranges from around ±0.5% p.a. for ages 55 to 80, to the order of ±1.0% p.a. for younger ages and for the oldest ages.

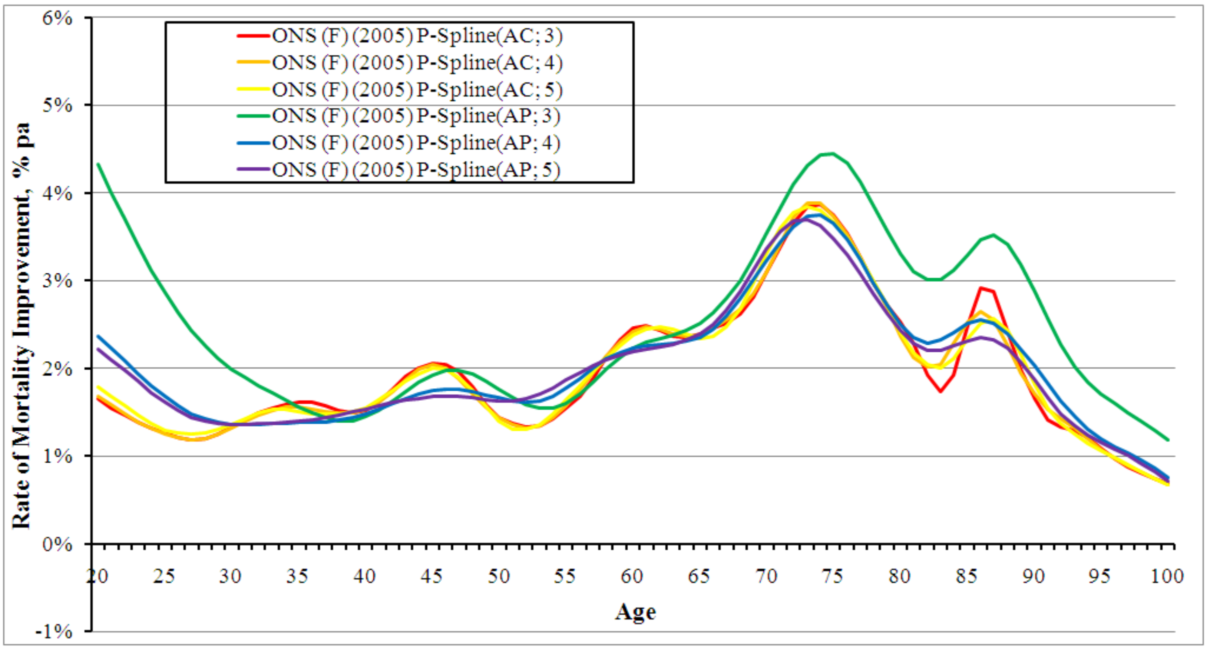

Figure 3.3: Annualised rates of mortality improvement in 2005: Comparison of estimates derived from alternative P-Spline models fitted to the ONS dataset for the female population of England &Wales

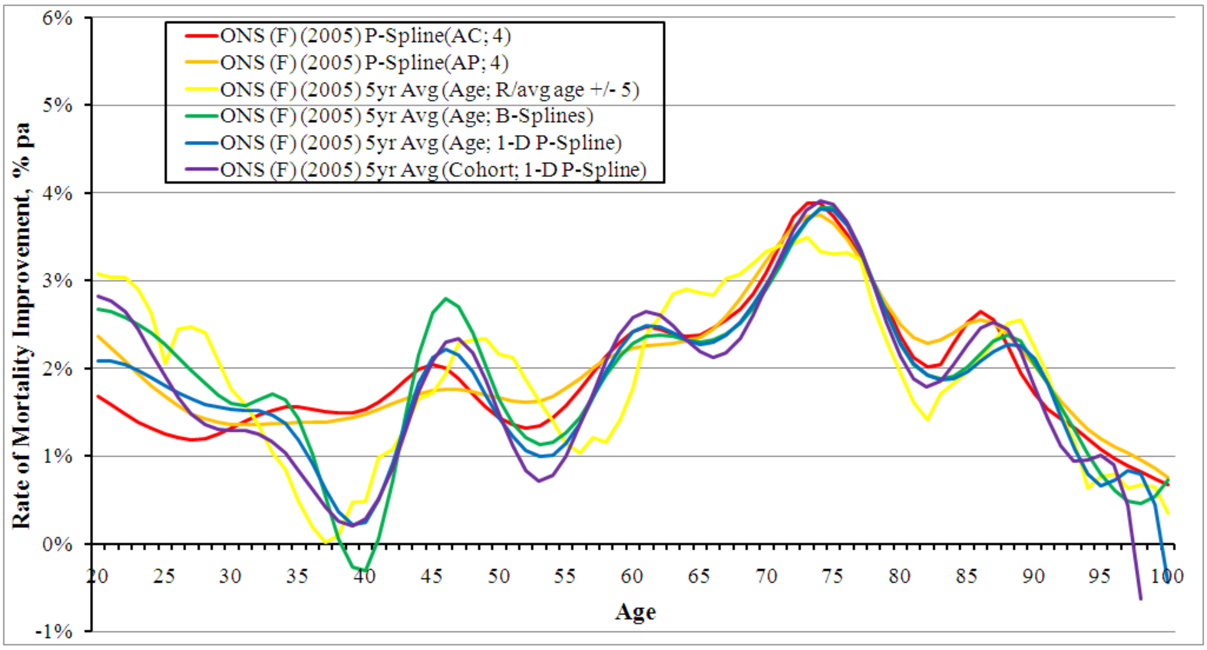

Figure 3.4: Annualised rates of mortality improvement in 2005: Comparison of estimates derived from fitted P-Spline models and various alternative methodologies applied to the ONS dataset for the female population of England &Wales

3.2.2. Insured and Pensioner Datasets

Having performed some analysis of current rates of mortality improvement observed for the population of England & Wales as a whole, it is natural to consider how experience may vary for subsets of the population. This section presents the results from an analysis of datasets covering Insured Lives and Pensioners.

The CMI datasets used for this analysis contain data up to 2006. Based on its findings for ONS data, the Working Party used 2004, i.e. again 2 years inside the edge of the available data, as the latest year for which sufficiently robust estimates of rates of mortality improvement can be made. Estimates of rates of mortality improvement in 2004 have been derived using the same set of methodologies as were applied to population data and the results (for male lives) are presented in the following charts.

Figure 3.5 shows estimates derived by applying P-Spline models to the CMI Permanent Assurances dataset for male lives assured.

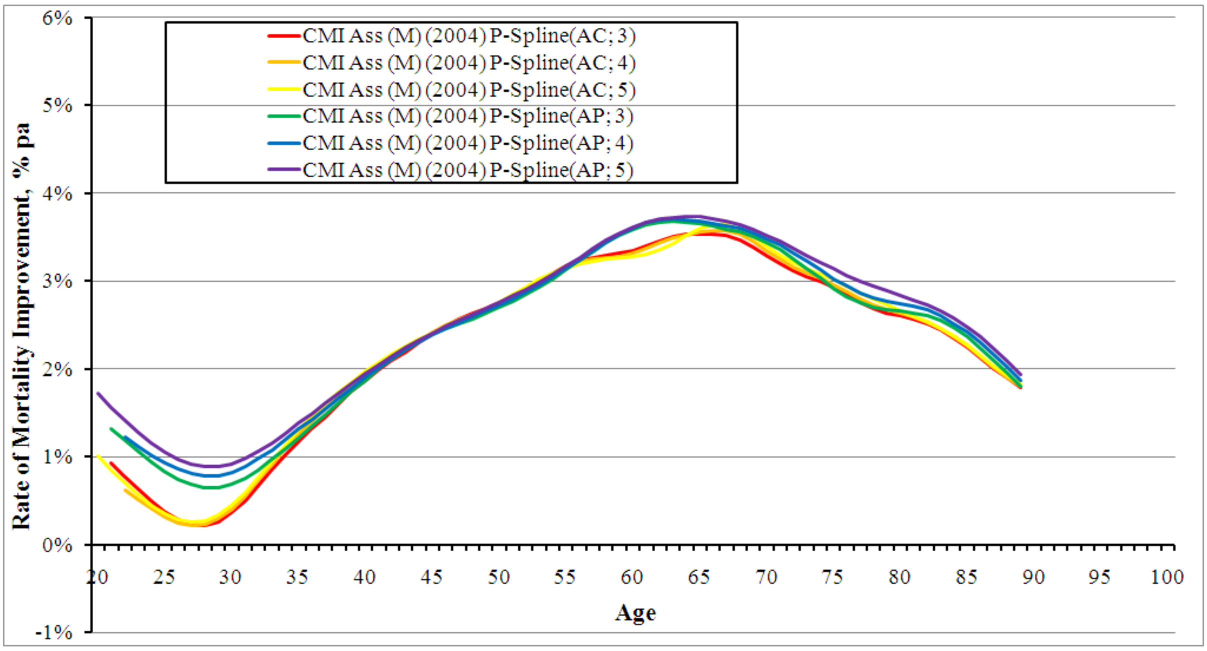

Figure 3.5: Annualised rates of mortality improvement in 2004: Comparison of estimates derived from alternative P-Spline models fitted to the CMI Permanent Assurances dataset for male lives assured

We observe that for this dataset (much smaller than the population datasets):

- The overall variation between alternative P-Spline models is second or third order compared to the variations by age in rates of mortality improvement.
- There are fewer features shown by the fitted models for CMI data than for population data, reflecting a greater degree of smoothing.
- There is more difference by model (age-cohort v age-period) than by knot-spacing for this dataset.

Figure 3.6 compares the rates of mortality improvement for 2004 derived from age-period and age-cohort P-Spline models with 4-year knot-spacing against rates derived using the selected alternative methodologies.

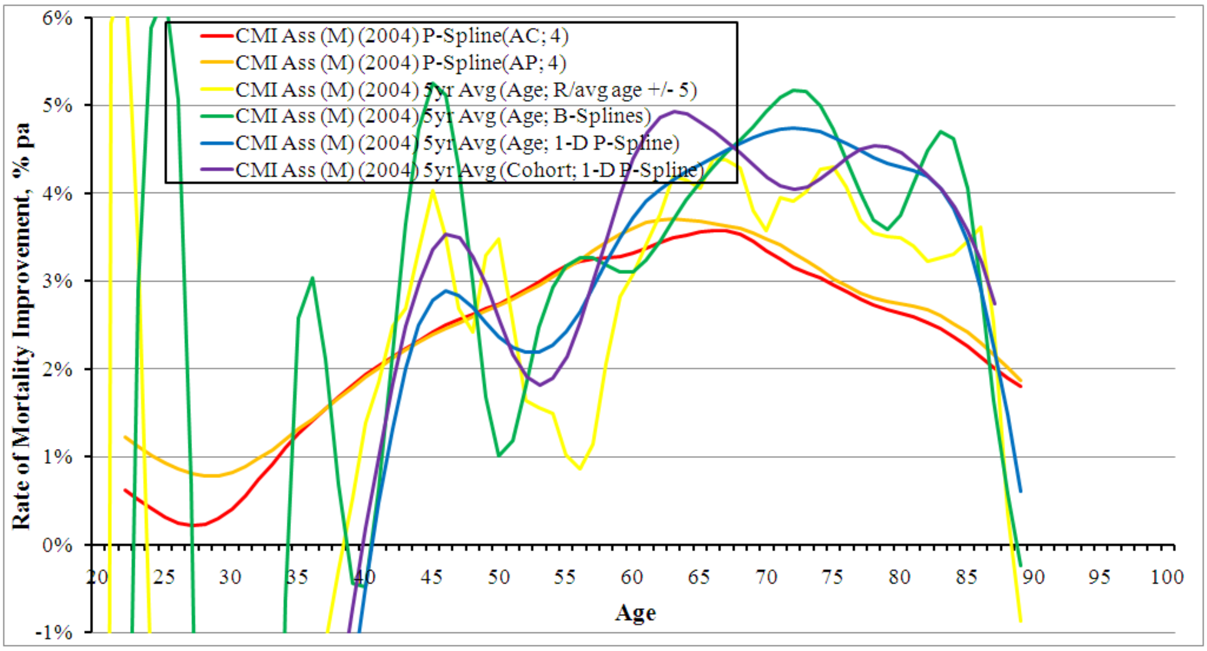

Figure 3.6: Annualised rates of mortality improvement in 2004: Comparison of estimates derived from fitted P-Spline models and various alternative methodologies applied to the CMI Permanent Assurances dataset for male lives assured

Figure 3.6 reveals a very different picture to that observed for the population datasets. For the CMI Permanent Assurances dataset:

- The correlation between estimates from P-Splines and from the alternative methodologies is much weaker than that observed for the population data.
- There is considerably greater estimation uncertainty – say ±1.0% p.a. at best, and rather more than that outside the age range 50 to 80 – than there is for population data.
- There is considerable difficulty in identifying true features of the underlying experience: the P-Spline models suggest a single peak in rates of mortality improvement around age 60 to 65, whereas the alternative measures hint at a peak around age 70 to 75 and perhaps at multiple local peaks.
- The alternative measures correlate better against each other than against the P-Spline estimates: the differences between the two groups reflect relatively heavy smoothing by age in the P-Spline models.
- Therefore the selection of P-Spline models versus alternative methodologies is more important in the context of estimating current rates of mortality improvement from CMI data than it is for population data.

Figures 3.7 and 3.8 show estimates of rates of mortality improvement derived by applying P-Spline models and the alternative methodologies to the CMI Life Office Pensioners dataset for male pensioners.

From Figure 3.7 we draw very similar observations and conclusions as for the CMI Permanent Assurances dataset in Figure 3.5. Figure 3.8 also shares many similarities with Figure 3.6 in terms of comparisons between measures and estimation errors. Again there is considerable difficulty in identifying true features of the underlying experience: the P-Spline models show a single peak in rates of mortality improvement around age 72, whereas the alternative measures suggest the peak is later – around age 75 – and somewhat higher.

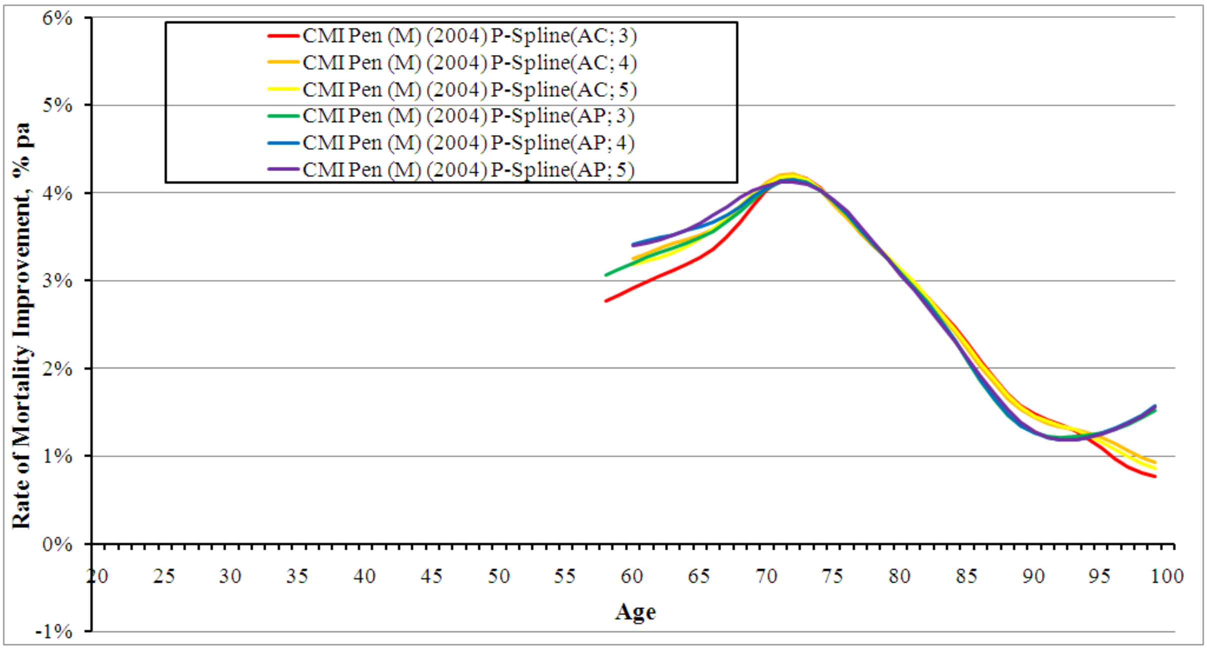

Figure 3.7: Annualised rates of mortality improvement in 2004:
Comparison of estimates derived from alternative P-Spline models fitted to the CMI Life Office Pensioners dataset for male pensioners

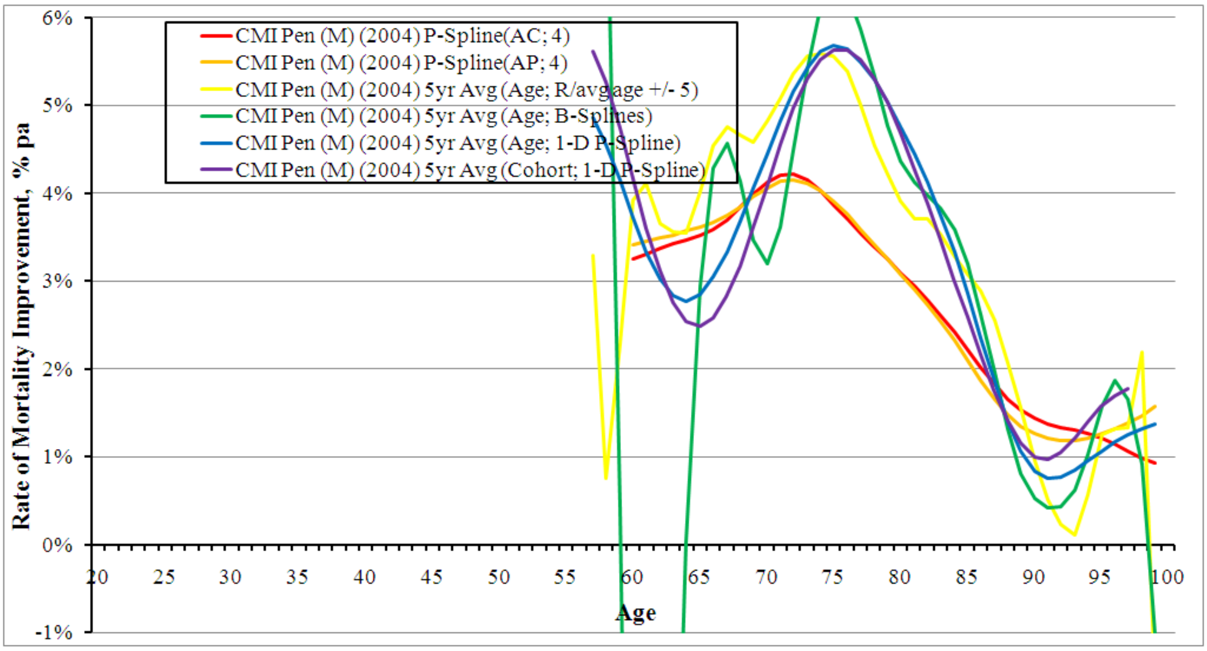

Figure 3.8: Annualised rates of mortality improvement in 2004:
Comparison of estimates derived from fitted P-Spline models and various alternative methodologies applied to the CMI Life Office Pensioners dataset for male pensioners

Selected estimates of the rates of mortality improvement for the population of England & Wales, the two CMI datasets, and the Club Vita dataset are compared in Figure 3.9. The estimates shown are for 2004 and for male lives only.

The estimates using ONS and CMI data have been derived using either a P-Spline model (an age-cohort model with 4-year knot spacing), or the 1-D P-Spline methodology (5-year averaging by age of rates of mortality improvement based on mortality rates for 2001-2006 smoothed using 1-dimensional P-Spline models fitted to each individual year's data).

For Club Vita the P-Spline model results shown are those provided by Club Vita to the Working Party. The age-cohort P-Spline model used in these results had 5-year knot spacing by age and 4-year spacing by cohort, purely for illustration purposes. Results on the alternative smoothing methodologies (including different P-spline variants) were not sought from Club Vita.

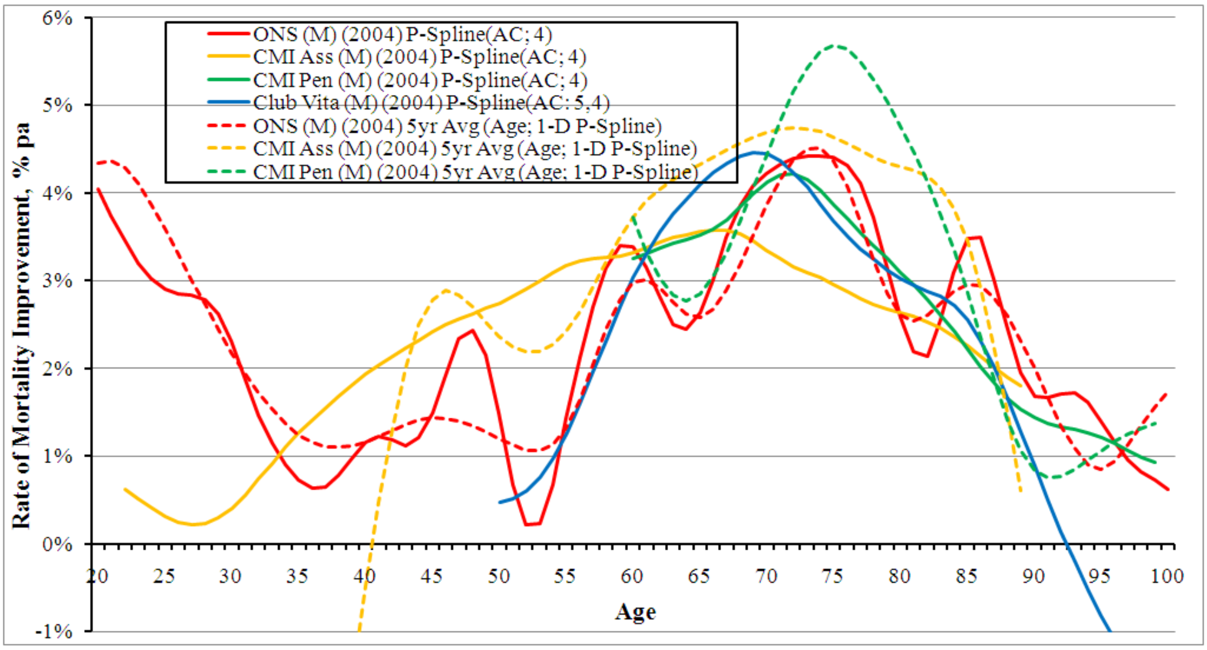

Figure 3.9: Annualised rates of mortality improvement in 2004: Comparison of sample estimates derived for male live using 4 datasets – Population of England &Wales (ONS), CMI Permanent Assurances, CMI Life Office Pensioners and Club Vita (pensioners)

Figure 3.9 shows the significant additional estimation error for CMI datasets compared to the population dataset. Given the relative data volumes, the Working Party expects the estimation error on the Club Vita data would be rather closer to that of the CMI datasets than the population.

As previously noted, the relative strength of data for the population dataset enables a greater number of features to be observed in the data with confidence. Estimation uncertainty makes it difficult to determine the extent to which the insured lives and pensioner datasets share these features or possess materially different features.

Concentrating on the largest / clearest observed peaks of mortality improvement rates:

- Population data shows a clear peak around year-of-birth 1931.
- CMI Permanent Assurances data currently shows a peak around years-of-birth 1940-45
- CMI Life Office Pensioner data currently shows a peak around years-of-birth 1929-33
- Club Vita data currently shows a peak around year-of-birth 1935.

It is difficult to see any pattern in these estimated rates of mortality improvement to link peak-cohort year-of-birth to socio-economic group. In order to investigate the issues further, the next section considers a 2-dimensional view of past trends in mortality experience, showing results over time as well as by age.

3.3. Trends in Rates of Mortality Improvement

3.3.1. Population of England & Wales

Trends and patterns in rates of mortality improvement over time and age may be examined using a ‘heat map’ presentation. As with the study of current rates of mortality improvement it is first necessary to apply some form of smoothing to the data, and then important to check to what extent any patterns observed may be influenced by the smoothing methodology.

All the charts included in this section show results from analysis using P-spline models to smooth the data. The Working Party satisfied itself that the features noted are genuine through parallel analyses of the data using the alternative methodologies set out in section 3.1.3, however these results have not been included in the paper.

The heat maps presented in this section use a colour spectrum to indicate the annual rate of mortality improvement for each 1-year time/age cell, with red at the high end of the scale, representing falls in mortality rates of between 4.5% and 5.5% over the year, through to purple at the low end, representing mortality change of between -4.5% and -3.5% (i.e. increases in mortality). Local peaks, by age, in mortality improvement rates are shown by black markers in the relevant cells.

Figure 3.10 shows annual rates of mortality improvement derived by applying an age-period P-Spline model to the ONS dataset for the population of England & Wales over the period 1961-2007 and age range 18 to 102.

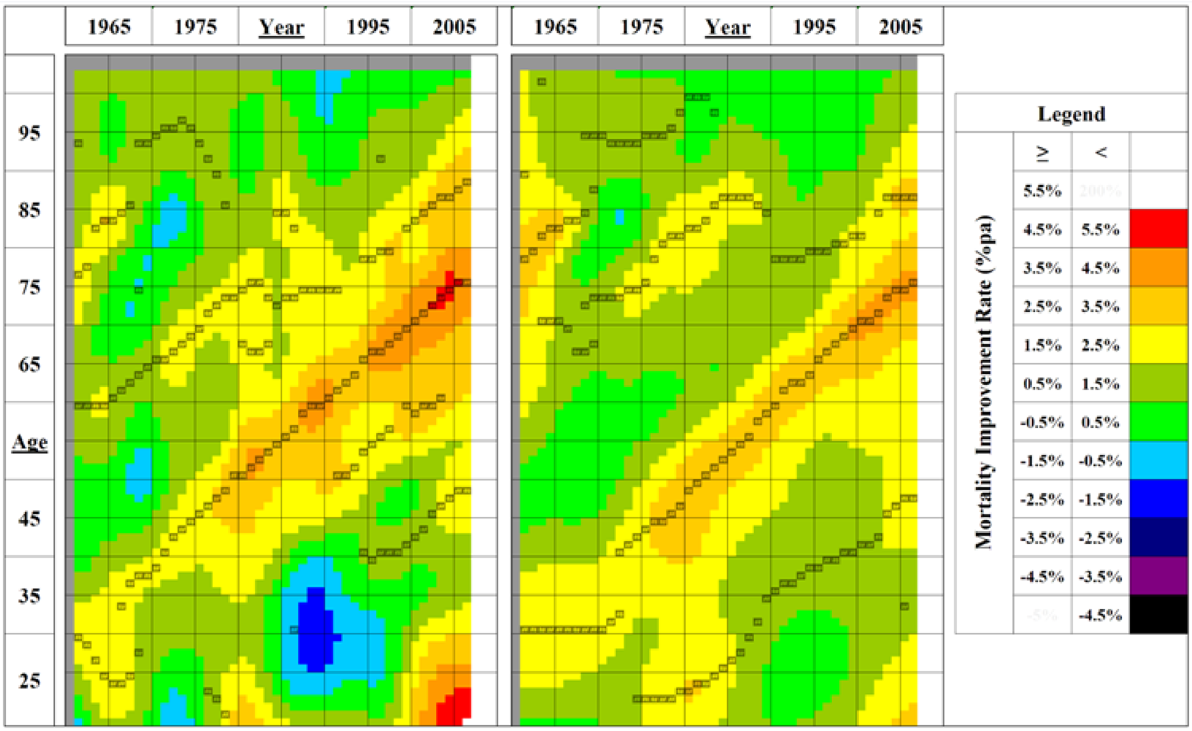

Figure 3.10: Heat map of estimated annual rates of mortality improvement: ONS dataset for the population of England &Wales Experience smoothed using an age-period P-Spline model (4-year knot-spacing) Left-hand panel: Males; Right-hand panel: Females

A number of features are clearly discernible. The warmer colours (yellows to reds), representing rapid rates of mortality improvement highlight the clear diagonal feature we recognise as the 'cohort effect' centred on the year-of-birth 1931 for both males and females. For males, there is also some evidence in support of cohort features around the years-of-birth 1943 and 1957.

These heat maps, for both males and females, also show a general steady increase in rates of mortality improvement across a wide range of ages, over the past 30 or so years.

Figure 3.11 shows annual rates of mortality improvement derived by applying an age-cohort P-Spline model to the same data.

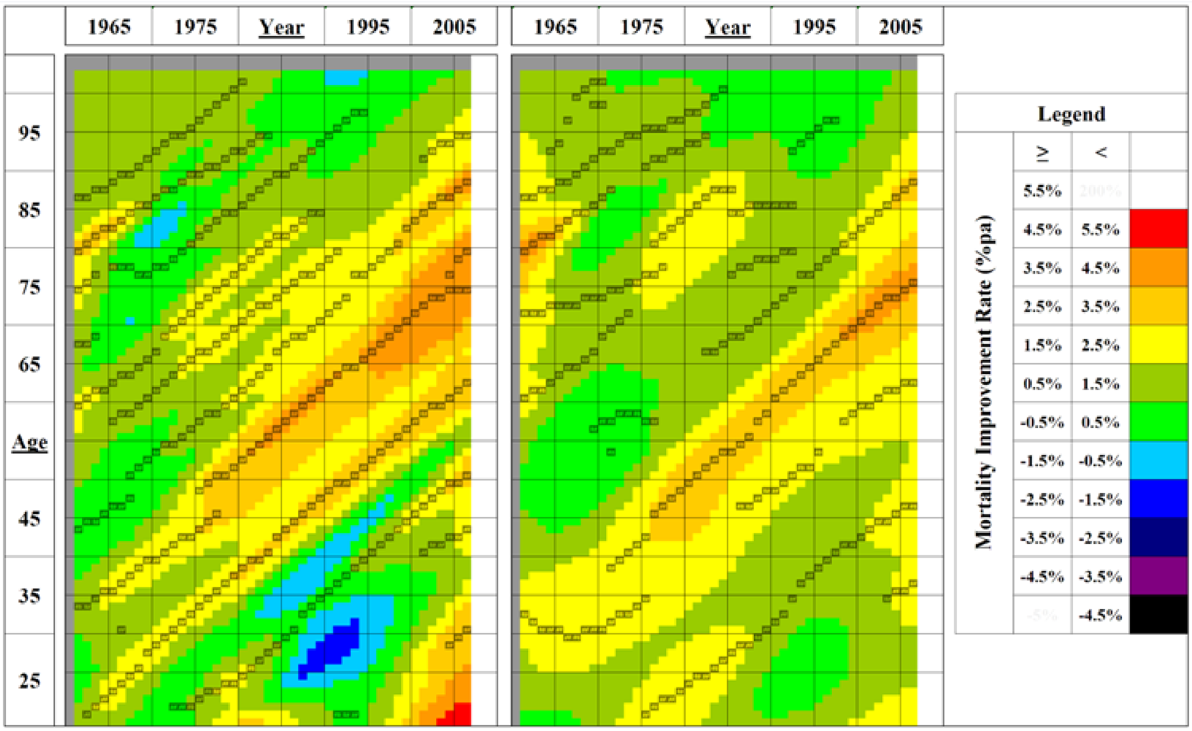

Figure 3.11: Heat map of estimated annual rates of mortality improvement: ONS dataset for the population of England &Wales Experience smoothed using an age-cohort P-Spline model (4-year knot-spacing) Left-hand panel: Males; Right-hand panel: Females

As expected, the age-cohort P-Spline model shows cohort features in greater number and strength than the age-period model (compared to Figure 3.10). Again, a very clear 'cohort effect' is seen for the year-of-birth around 1931 for both males and females, as well as the general, steady increase in rates of improvement over the last 30 or so years. Some evidence of cohort features for years-of-birth 1943 and 1957 is now seen for both genders, although the dataset for males offers stronger support for these features than the dataset for females.

The heat map analysis revealed a general, steady increase in rates of mortality improvement over the last 30 or so years. The feature is further investigated in Figure 3.12 which shows annual rates of mortality improvement, derived by applying an age-period P-Spline model, by year-of-birth cohort, for a series of calendar years covering 1970 to 2005 in 7-year steps.

The calendar-year lines are broadly parallel across a wide range of birth years, indicating strong and persistent cohort features (both peaks and troughs). In addition, these calendar-year lines also move steadily up the vertical axis, indicating across the board increases in rates of improvement. This suggests rates of population mortality change over the period could reasonably be decomposed into time and cohort features.

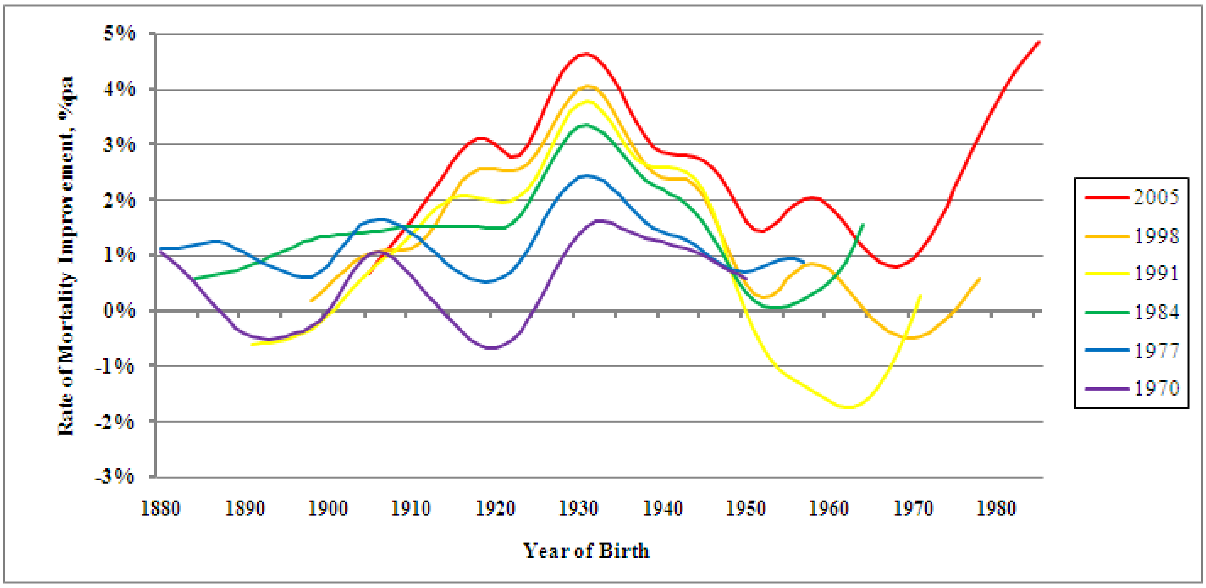

Figure 3.12: Estimated annual rates of mortality improvement, by cohort and year ONS dataset for the male population of England &Wales Experience smoothed using an age-period P-Spline model (4-year knot-spacing)

It is also notable that Figure 3.12 offers no evidence to suggest that rates of mortality improvement will immediately begin to tail off (as would be modelled by a straight line convergence to an assumed lower long-term rate).

Drawing from Figures 3.10, 3.11 and 3.12 together, there is clear evidence from the England & Wales population data of persistent year-of-birth cohort features, along with a more general, steady and sustained, increasing trend in rates of mortality improvement across a wide range of ages. These conclusions drive the proposal to separate ‘by age’ and ‘by cohort’ components of mortality change in the Model.

Before moving on from the population dataset, it is worthwhile to examine the pattern of mortality improvement rates at high ages.

The England & Wales population dataset used by the Working Party contained individual age data up to age 104. However, the methodology generally used by the ONS to estimate mid-year populations (taken as the exposure measure for this analysis) is only applied up to age 89. A different methodology is applied for ages 90 and above and this could lead to a discontinuity in derived mortality rates by age. The Working Party investigated this issue further and is satisfied that the change of methodology around age 90 does not cause a material distortion for the smoothed rates of mortality improvement it has calculated.

Figure 3.13 shows annual rates of mortality improvement, derived by applying an age-period P-Spline model, by age for a series of calendar years, this time covering 1985 to 2005 in 4-year steps.

As rates of change are shown here by age, rather than by year-of-birth, the movement of local peaks to the right with each calendar-year step illustrates cohort features; general increases in the rate of mortality improvement over time are shown by the peaks moving up the vertical axis.

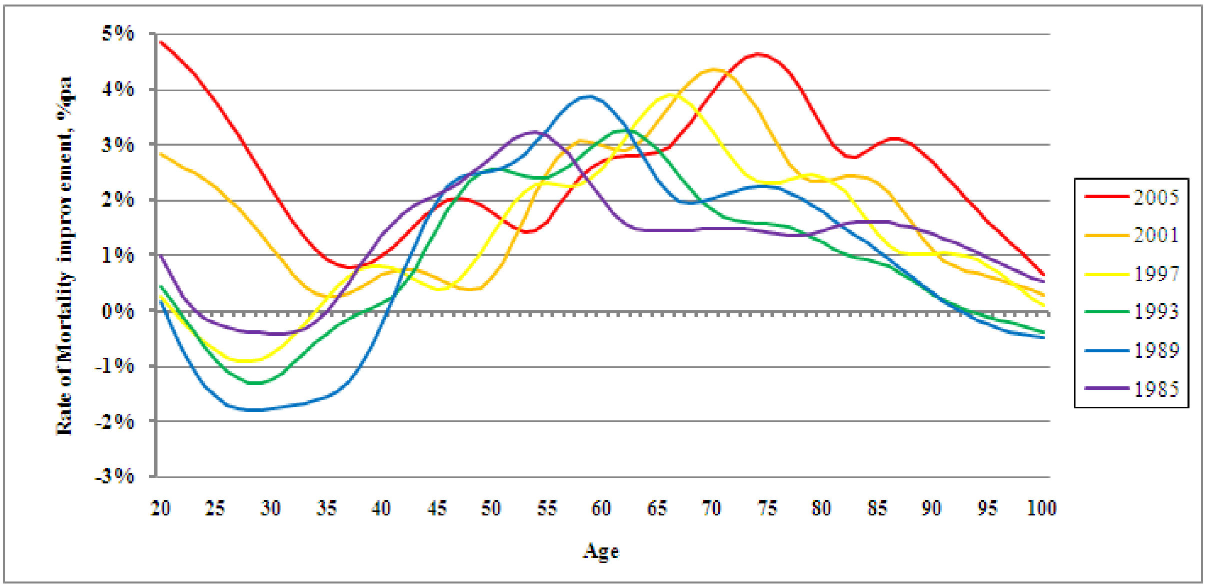

Figure 3.13: Estimated annual rates of mortality improvement, by age and year ONS dataset for the male population of England &Wales experience smoothed using an age-period P-Spline model (4-year knot-spacing)

Figure 3.13 shows that rates of mortality improvement seem to run down towards zero at high ages, say between 90 and 105. However, the estimated improvement rates at high ages display a near-parallel shift over the past two decades, so that the age-point of reaching around 'zero', that is stable mortality rates, appears to have been increasing over time.

3.3.2. Insured and Pensioner Datasets

Turning now to analysis of datasets covering Insured Lives and Pensioners, the CMI Permanent Assurances, CMI Life Office Pensioners and Club Vita (self-administered pension schemes) datasets are examined to see the extent to which there are common features, and how experience may vary for subsets of the population. Again, the Working Party used P-spline models and the alternative methodologies set out in section 3.1.3 in its analysis, but only the P-spline results are shown in the paper.

Figures 3.14, 3.15 and 3.16 show heat maps of annual rates of mortality improvement derived by fitting P-Spline models to the various datasets. In all three cases, these charts show fewer features than the equivalent heat maps for population data, reflecting the much lower data volumes for these specific population subgroups.

Figure 3.14 shows the estimates for CMI Permanent Assurances dataset. The data cover the age range 20 to 90, and the periods 1961-2006 for male lives and 1983-2006 for female lives.

The relatively low data volumes make it difficult to draw conclusions at all from the right-hand panel for females, and also for males at older ages, as data volumes fall away sharply above age 65 for this dataset.

The map for males hints a feature initially broadly in line with the population 1931-cohort feature, but shows a falling away, with the peak recently holding level on age as calendar years pass. This means that with the addition of more recent data, the picture is now very different to that leading to the CMI Working Paper 1 conclusion in respect of a peak in rates of mortality improvement for the 1926-cohort.

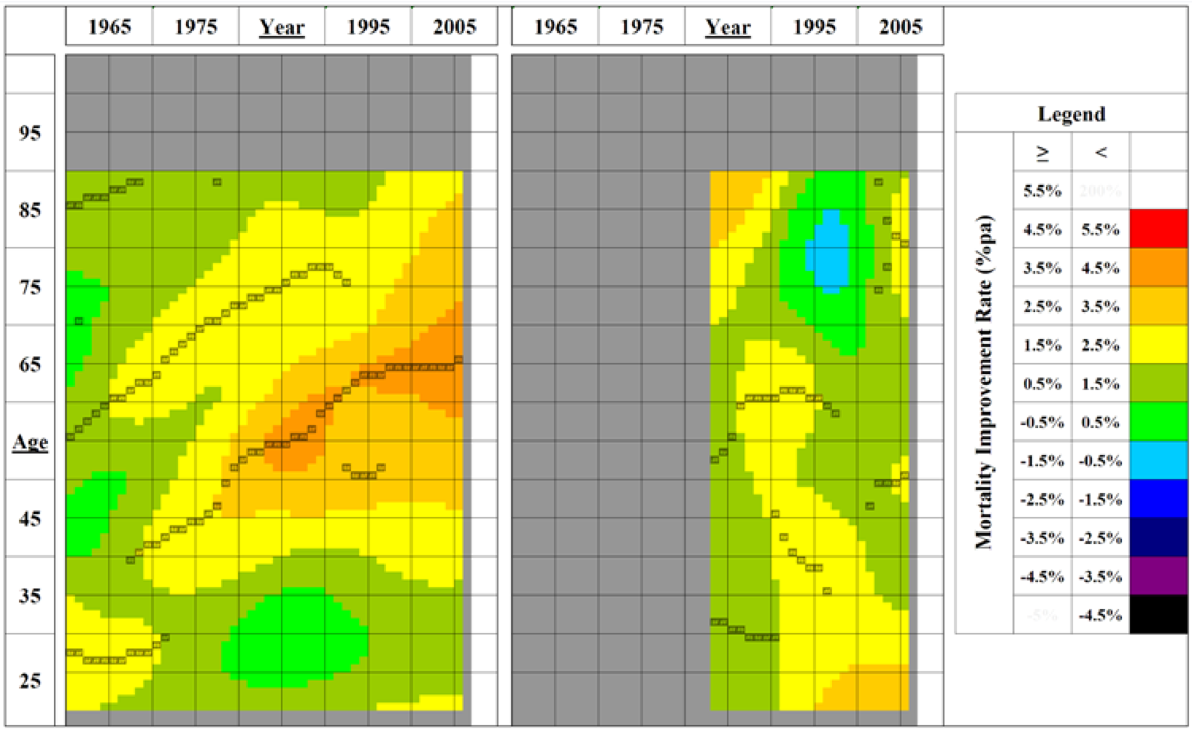

Figure 3.14: Heat map of estimated annual rates of mortality improvement: CMI Permanent Assurances dataset
Experience smoothed using an age-period P-Spline model (5-year knot-spacing)
Left-hand panel: Males; Right-hand panel: Females.

Whilst the results for the 1970s and 1980s support a peak for the 1930-cohort (not 1926), experience in the 1990s and beyond has seen an ‘age effect’ come to prominence, such that the peak falls away from the cohort progression (and was around year-of-birth 1940 by 2006).

Although the cohort features, initially at least, appear different to those seen in population data, Figure 3.14 does also show a general, steady increase in rates of mortality improvement rates, across a wide range of ages, over the past 30 years, which strongly echoes that similar feature in population data.

Figure 3.15 shows the estimates for CMI Life Office Pensioners dataset. The available data cover the age range 60 to 100, and the period 1983-2006. The data volumes for this dataset are concentrated above age 65, in sharp contrast to the distribution for the Permanent Assurances dataset. Even so, it is still difficult to draw conclusions from the right-hand panel for females.

Once again, the heat map for males confirms a general increasing trend in rates of improvement over recent times, and shows the peak of rates of mortality improvement falling below the initial cohort-based line. Whilst the results to the mid-1990s perhaps support a peak for the 1928-cohort (not 1926), subsequent experience has seen an 'age effect' beginning to dominate, such that the peak falls away from the cohort progression (and was around year-of-birth 1933 by 2006).

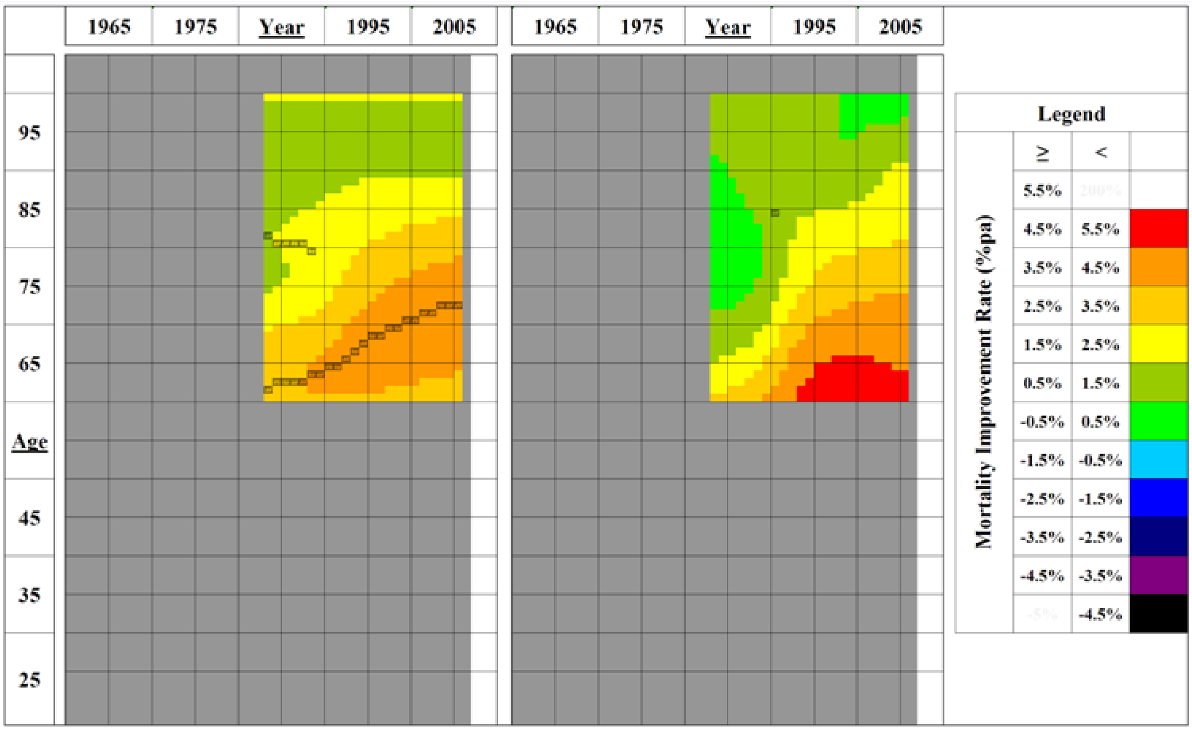

Figure 3.15: Heat map of estimated annual rates of mortality improvement: CMI Life Office Pensioners dataset
Experience smoothed using an age-period P-Spline model (5-year knot-spacing)
Left-hand panel: Males; Right-hand panel: Females.

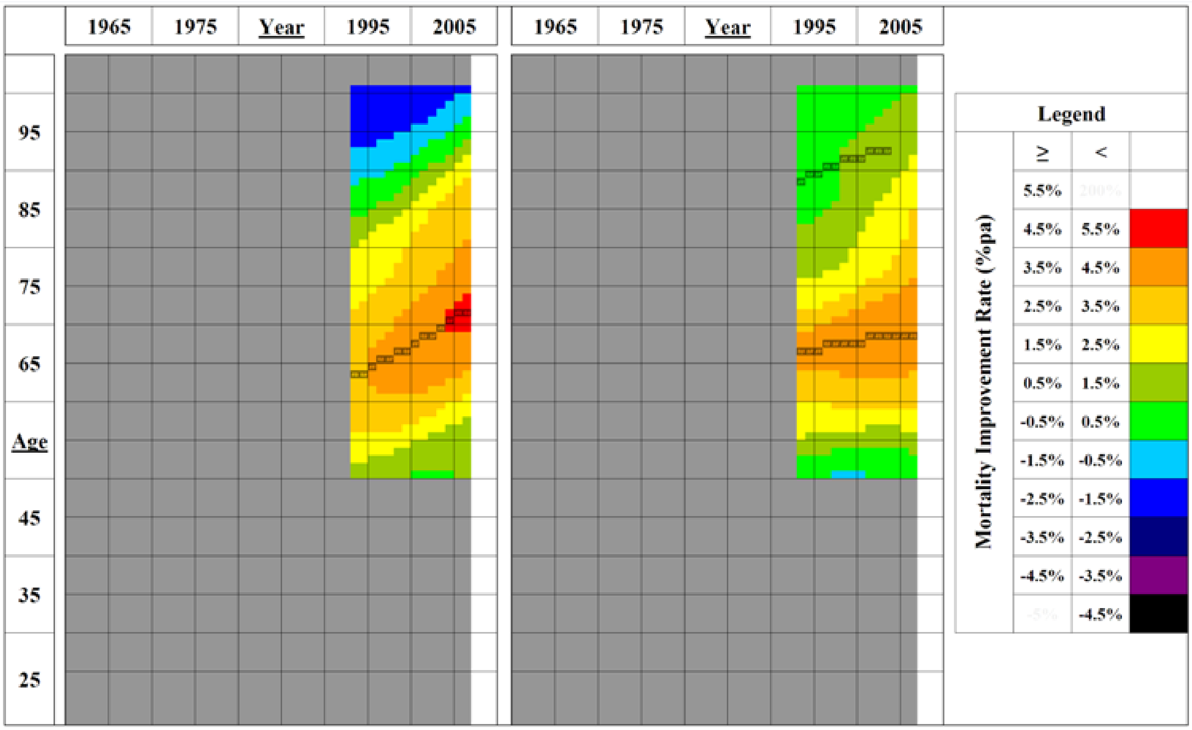

Figure 3.16: Heat map of estimated annual rates of mortality improvement: Club Vita dataset (self-administered pension schemes) Experience smoothed using an age-cohort P-Spline model (5/4-year knot-spacing) Left-hand panel: Males; Right-hand panel: Females.

Figure 3.16 shows the estimates for the Club Vita dataset. The Working Party was supplied with the (P-Spline) smoothed rates shown in the chart, covering the age range 50 to 100, and the period 1994-2007. Whilst the data volume is somewhat larger than the CMI Pensioner dataset for males and substantially larger for females, the analysis period is significantly shorter.

The heat map for males shows reasonable support for a peak in rates of mortality improvement for the 1935 year-of-birth cohort. However, there is also evidence, similar to that from the CMI datasets, of the peak drifting below the initial cohort-based line. For females, the map appears to show an age feature rather than a cohort feature. Once again the maps also show a general, steady increase in rates of mortality improvement, across a wide range of ages, over the period of data.

Figure 3.17 shows annual rates of mortality improvement, derived by applying an age-period P-Spline model, for the CMI Permanent Assurances dataset for male lives assured. Rates are shown by year-of-birth cohort for a series of calendar years, covering 1970 to 2004 in 7-year steps. The chart is directly comparable (as far as data availability allows) with Figure 3.12 (analysis of the dataset for the male population of England & Wales).

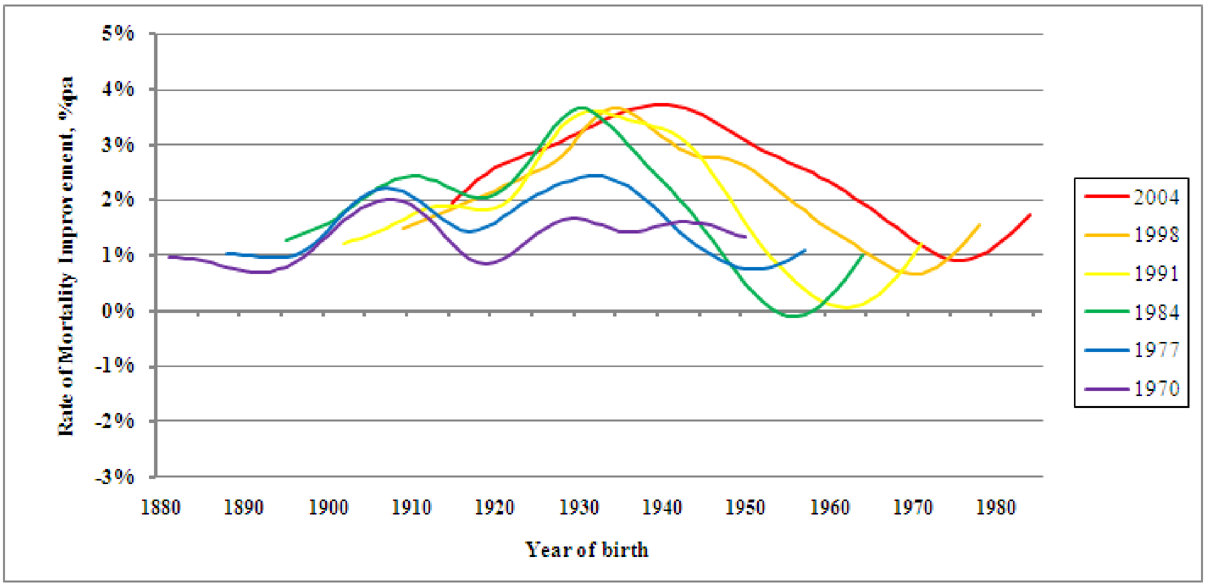

Figure 3.17: Estimated annual rates of mortality improvement, by cohort and year CMI Permanent Assurances dataset for male lives assured Experience smoothed using an age-period P-Spline model (5-year knot-spacing)

Figure 3.17 shows features similar to those seen for population data, namely strong cohort features and a general increase in rates of mortality improvement over time. However, there is a significant shift in the pattern after 1991, with the main peak shifting from year-of-birth 1931 or so to around year-of-birth 1939.

It is also notable that the lines representing data for more recent years show fewer features than the lines for the 1970s and 1980s. The data volumes in this dataset have fallen significantly over the period shown, so that the latest data requires smoothing over broader age ranges and only supports evidence of more general features.

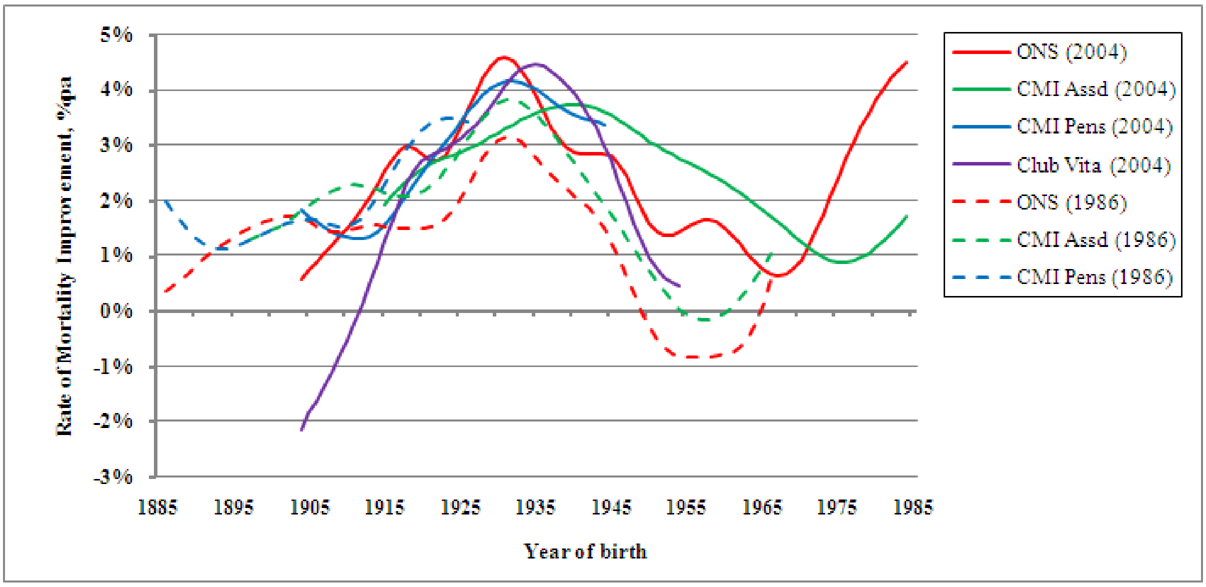

Figure 3.18: Estimated annual rates of mortality improvement: by cohort, for 1986 and 2004, for Population, CMI and Club Vita datasets, male lives Experience smoothed using P-Spline models

Figure 3.18 shows annual rates of mortality improvement, derived by applying age-period P-Spline models, for the population of England & Wales, CMI Permanent Assurances and CMI Life Office Pensioner datasets. The rates are shown only for male lives, and plotted by year-of-birth cohort for two calendar years, 1986 and 2004, to illustrate the shifts which have occurred over the intervening period. Rates derived from the Club Vita dataset (using an age-cohort P-Spline model) are also shown for 2004, for comparison.

The main peak in rates of mortality improvement by cohort appears to be reasonably stable for male lives in the population dataset, although the absolute level of rates of change is generally higher in 2004 than it was in 1986. Analysis of data within the period 1986 to 2004 shows that these statements also broadly hold for the CMI Life Office Pensioner dataset. However, an apparently significant shift across year-of-birth cohorts is again most visible for the CMI Permanent Assurances dataset. Interestingly, further analysis of the data, using the alternative smoothing methodologies set out in section 3.1.3, suggests that the shift in the peaks by cohort does not appear to affect the relative peak for years of birth 1955-1960.

The Insured and Pensioner datasets all show some evidence of cohort-type features but with a drifting of the peak across year-of-birth cohorts over the last decade or so. It is possible that this pattern results from the combination of genuine cohort features with the observed, more general, trend of increasing rates of mortality improvement. This is explored further in the next section.

3.3.3. Observations on the Effect of Data Volumes

In analysing mortality data for rates of change, for example by fitting a P-Spline mortality surface, a balance must be struck between goodness-of-fit and smoothness. The degree of smoothing applied tends to increase as data volumes reduce and identification of features in the data becomes increasingly uncertain (in a statistical sense).

From the analysis presented in section 3.2 and earlier in this section **Error! Reference source not found.**, apparently different features have been observed when comparing the population dataset against the Insured and Pensioner datasets. The Working Party wished to explore the extent to which the features noted in Insured and Pensioner datasets might actually be consistent with the more refined features observed in the population dataset but hidden under a greater degree of smoothing.

For this further investigation, age-period P-Spline surfaces were fitted to data based on the mortality rates observed in the male population of England & Wales. Figures 3.19 and 3.20 show the results of this analysis.

In each case the left-hand panel shows the heat map derived by fitting the P-Spline model to the actual data – that is, the trade-off between goodness-of-fit and smoothness is based on population-level data volumes.

The right-hand panels show the heat maps derived by fitting a P-Spline model to a scaled-down version of the population dataset. For each 1-year age-year cell, both the deaths and exposure data have been reduced by applying the ratio of CMI data to population data exposures for the cell. This produces a surface which reflects the changes in population mortality rates, by age and over time, but with the trade-off between goodness-of-fit and smoothness based on CMI dataset volumes. Figure 3.19 uses data volumes from the CMI Permanent Assurances dataset, and Figure 3.20 those from the CMI Life Office Pensioners dataset. In each case the age and calendar year ranges, for both left and right panels, were set to cover the overlap of the relevant population and CMI data.

Comparing the left and right-hand panels in Figures 3.19 and 3.20, the additional smoothing in the P-Spline surface resulting from the scaled-down data is clear. Although the underlying rates of mortality are the same for both panels, the effect of notionally scaling down the data volumes is that the right-hand panel loses many of the features shown in the (actual data volume) analysis in the left-hand panel.

Further, the 1931 year-of-birth 'cohort effect' so clearly observed as a strong diagonal feature in the analysis of population data (left-hand panels) appears to drift below the initial cohort-based line in the analyses scaled to CMI data volumes (right-hand panels). The change of scale of data creates a picture of population mortality change which shows features similar to those seen in the analysis of Insured and Pensioner datasets.

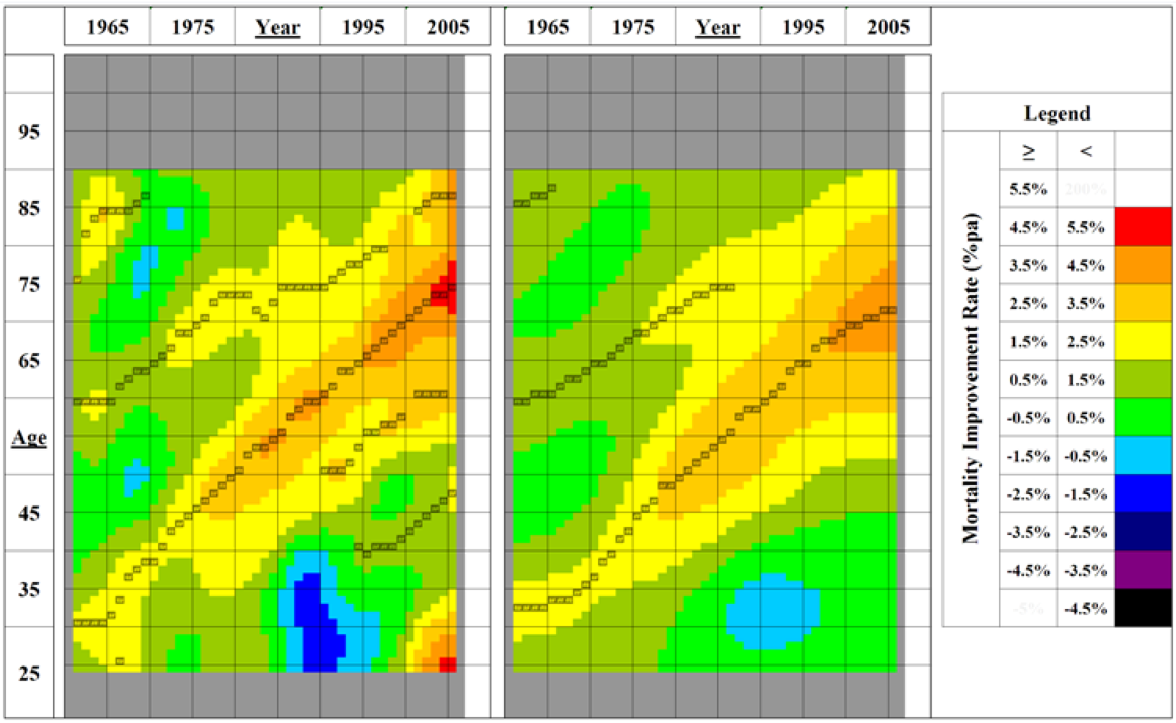

Figure 3.19: Heat map of estimated annual rates of mortality improvement: ONS dataset for the male population of England &Wales experience smoothed using age-period P-Spline models (5-year knot-spacing)
Left-hand panel: P-Spline applied to actual (population) data volumes
Right-hand panel: P-Spline applied to data scaled to CMI Permanent Assurances.

This work therefore raises questions over the strength of the initial observations previously drawn using heat maps to compare patterns in rates of mortality improvement for the population and Insured and Pensioner datasets. For example, for the CMI Permanent Assurances dataset, Figure 3.14 appeared to show an initial 1930-cohort feature evolving such that the peak of improvement rates fell away from the cohort progression and drifted slowly across year-of-birth cohorts over the last decade or so. However, this pattern is consistent with an amalgamation of cohort and age/period effects similar to those seen in population data. Whilst at population level there is sufficient strength of data to separate these two concurrent features, this analysis shows that the lower data volumes of the Insured and Pensioner datasets are insufficient to allow such component features to be separately identified.

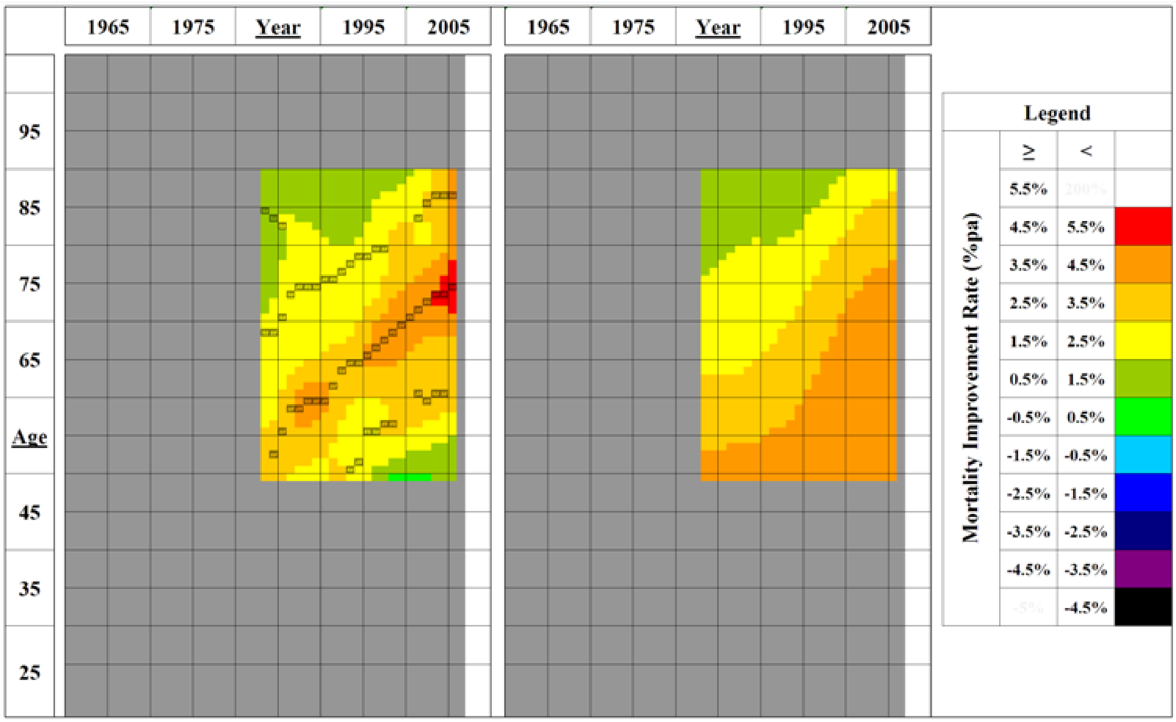

Figure 3.20: Heat map of estimated annual rates of mortality improvement: ONS dataset for the male population of England &Wales
Experience smoothed using age-period P-Spline models (5-year knot-spacing)
Left-hand panel: P-Spline applied to actual (population) data volumes
Right-hand panel: P-Spline applied to data scaled to CMI Life Office Pensioners.

3.3.4. Trends in Mortality by Social Class

To further investigate patterns of change for different sub-groups of the population, the Working Party analysed rates of mortality improvement by social class using ONS data.

In 2007 the ONS published age-specific mortality rates by social class for the period 1972-2005. The mortality rates were derived by the ONS using data from the long-running Longitudinal Study, which has monitored 1% of the population of England & Wales over time. Social class for each individual is taken as at their date of entry to the study based on the following classification:

- Class I: Professionals
- Class II: Managerial and technical intermediate
- Class IIIN: Skilled non-manual
- Class IIIM: Skilled manual
- Class IV: Partly skilled
- Class V: Unskilled

Figure 3.21 shows the average annual rate of mortality improvement for males by age group and social class between the period 1972-76 and the period 2002-05. In order to see the improvement patterns more clearly mortality rates were calculated for three distinct groups: Classes I, II and IIIN combined, Class IIIM, and Classes IV and V combined. Aggregate rates of improvement were also calculated for all social classes.

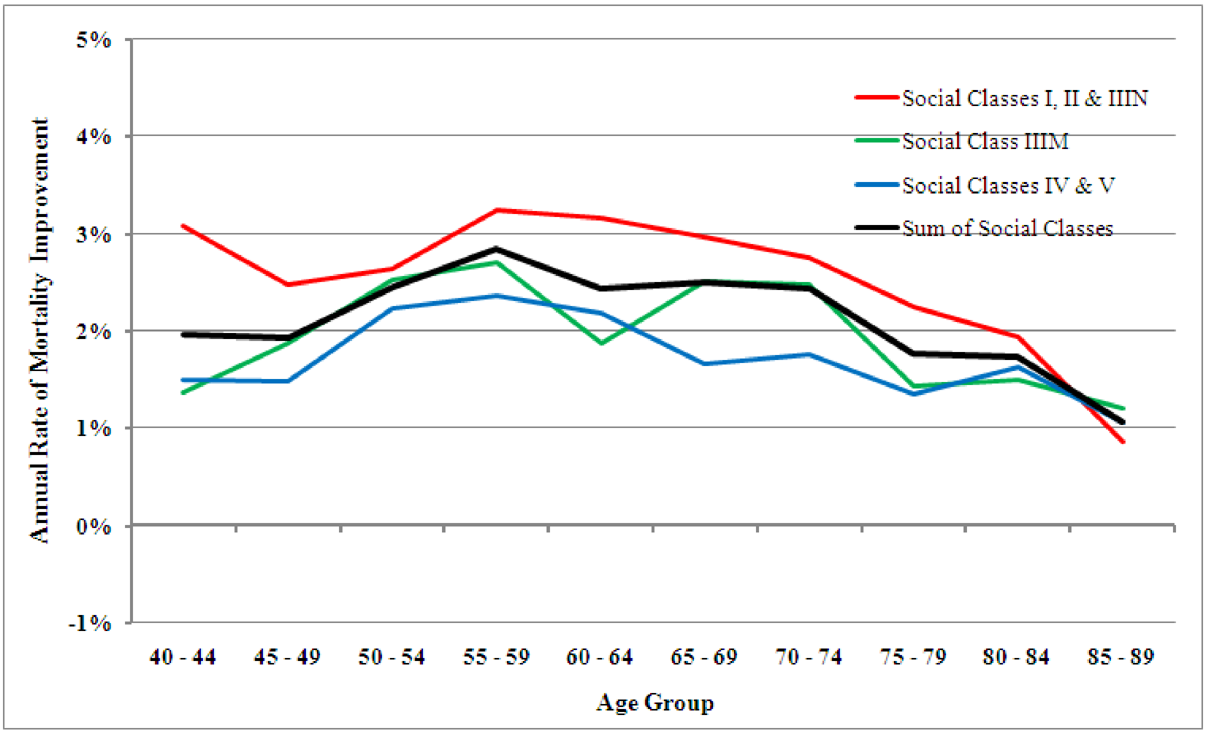

Figure 3.21: Average annual rate of mortality improvement for males in England & Wales, 1972-76 to 2002-05, by age group and social class

Figure 3.21 shows a clear differential in the average rate of improvement, with mortality rates falling more rapidly for the more affluent social class groups. The differentials in improvement rate (in absolute terms) have been relatively constant up to the 75-79 age group, but are less marked for ages 80-84 and not present at all for those aged 85-89.

It should be borne in mind that those in poor health at the beginning of the study may have been more frequently allocated to lower social class groups, and this potential 'reverse selection' will have unwound over time. Moreover, as the social class used for each individual is set at their point of entry to the study it will not reflect any changes to social class (i.e. affluence or deprivation) throughout the individual's life. However, the relative differentials in rates of mortality improvement appear broadly consistent with observations based on CMI data and experience in other developed countries.

Figures 3.22 and 3.23 show the same data spilt into different periods, i.e. from 1972-76 to 1987-1991 and from 1987-1991 to 2002-05.

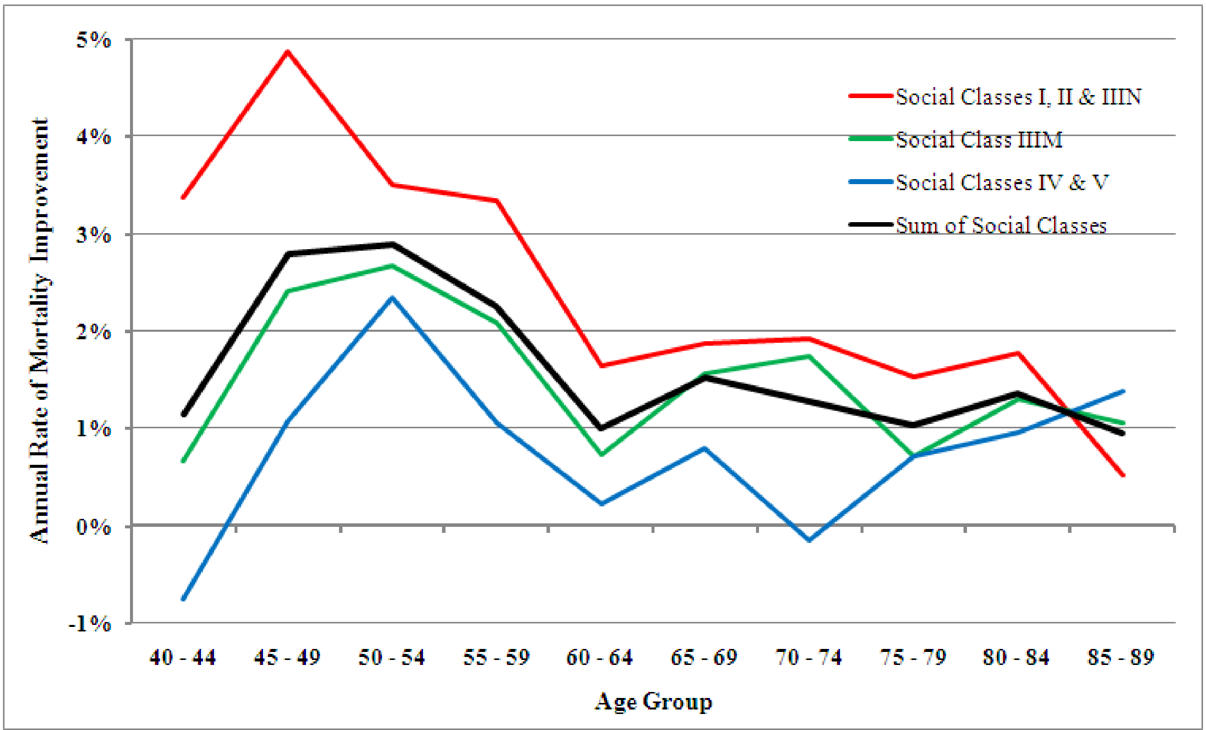

Figure 3.22: Average annual rate of mortality improvement for males in England & Wales, 1972-76 to 1987-91, by age group and social class

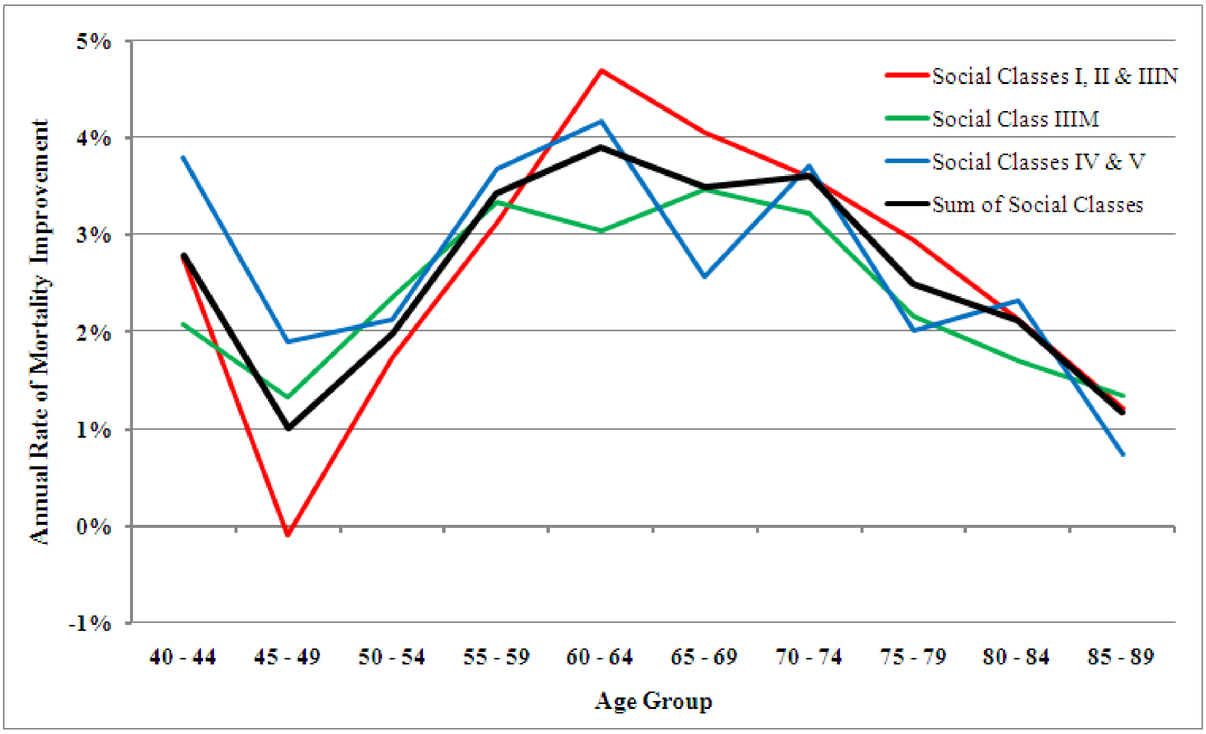

Figure 3.23: Average annual rate of mortality improvement for males in England & Wales, 1987-91 to 2002-05, by age group and social class

Mortality improvement rate differentials by social class are clearly shown by Figure 3.22 to have been significant over the period from 1972-76 to 1987-91. However, the improvements for the more recent period, illustrated by Figure 3.23, do not demonstrate such an obvious pattern. Indeed, most recently, average rates of improvement for age groups under 60 have been more rapid for the lowest social class groups (IV & V combined) than the highest groups (I, II & IIIN combined).

As noted above, the fact that social class is set at the point of entry to the study may have contributed to the pattern of change observed.

The social class mortality patterns and differentials are generally similar, but weaker, for females compared to males. Equivalent values to those given for males are shown in Figures 3.24, 3.25 and 3.26.

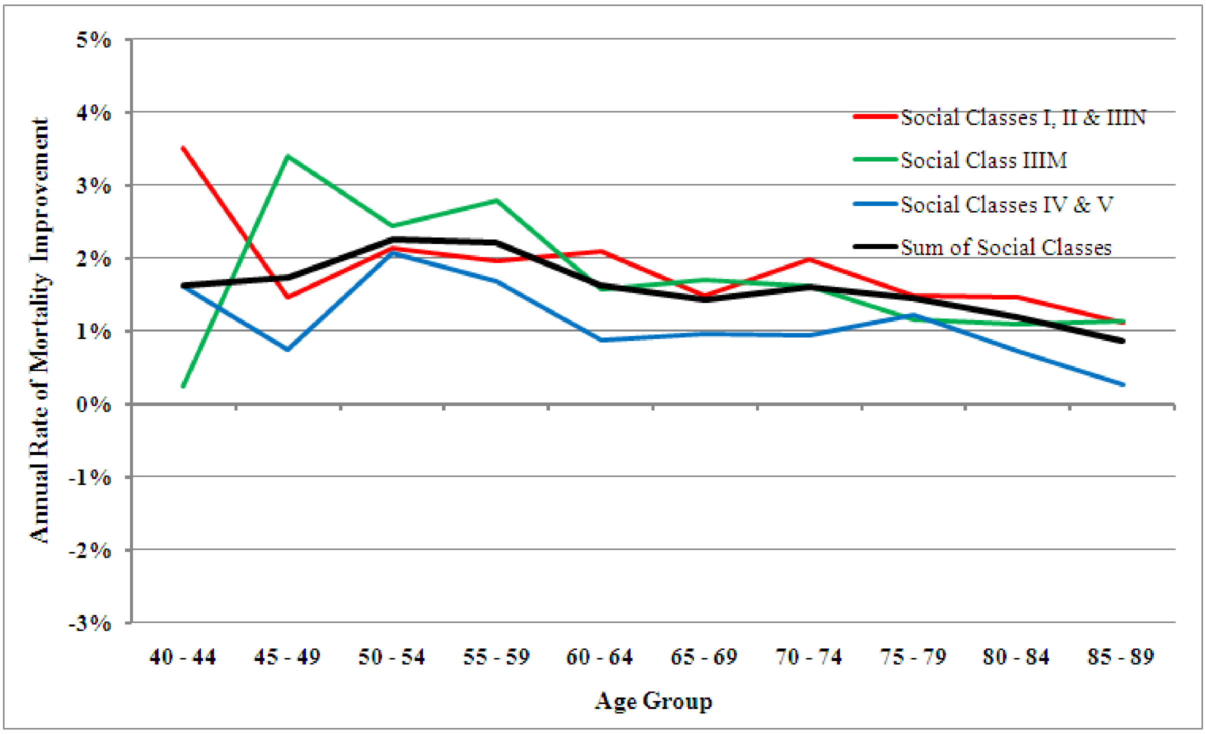

Figure 3.24: Average annual rate of mortality improvement for females in England & Wales, 1972-76 to 2002-05, by age group and social class

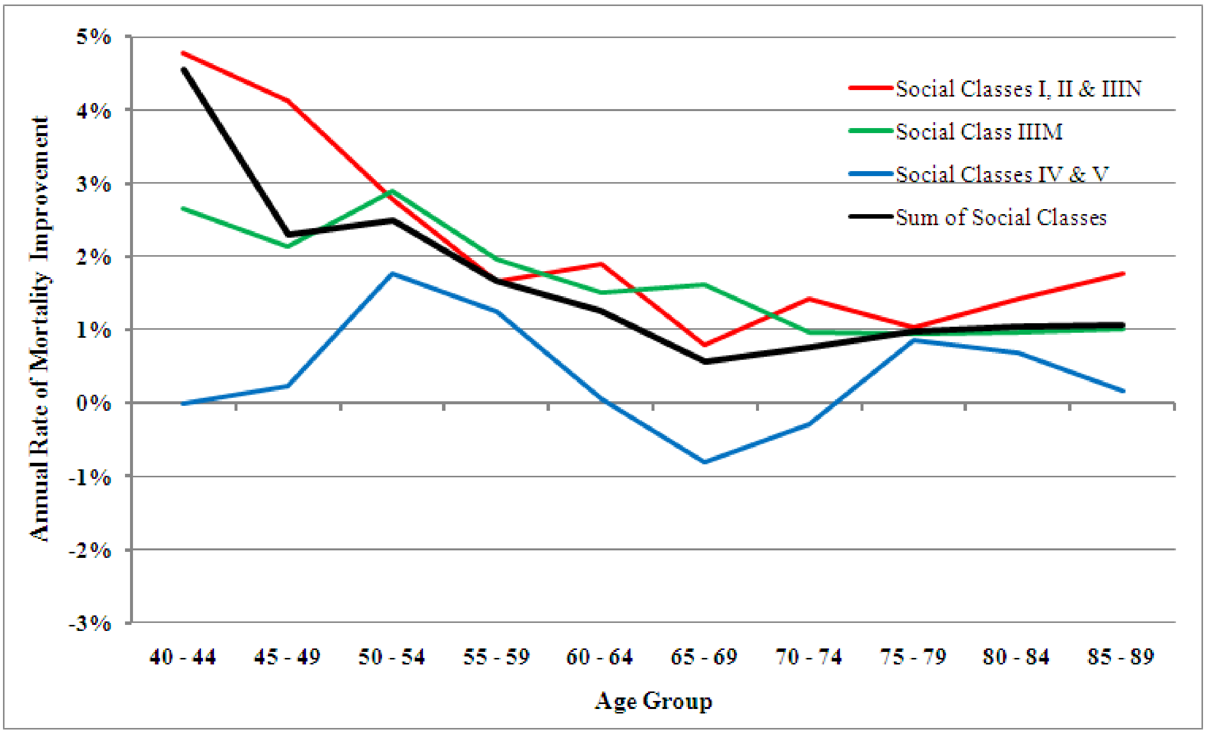

Figure 3.25: Average annual rate of mortality improvement for females in England & Wales, 1972-76 to 1987-91, by age group and social class

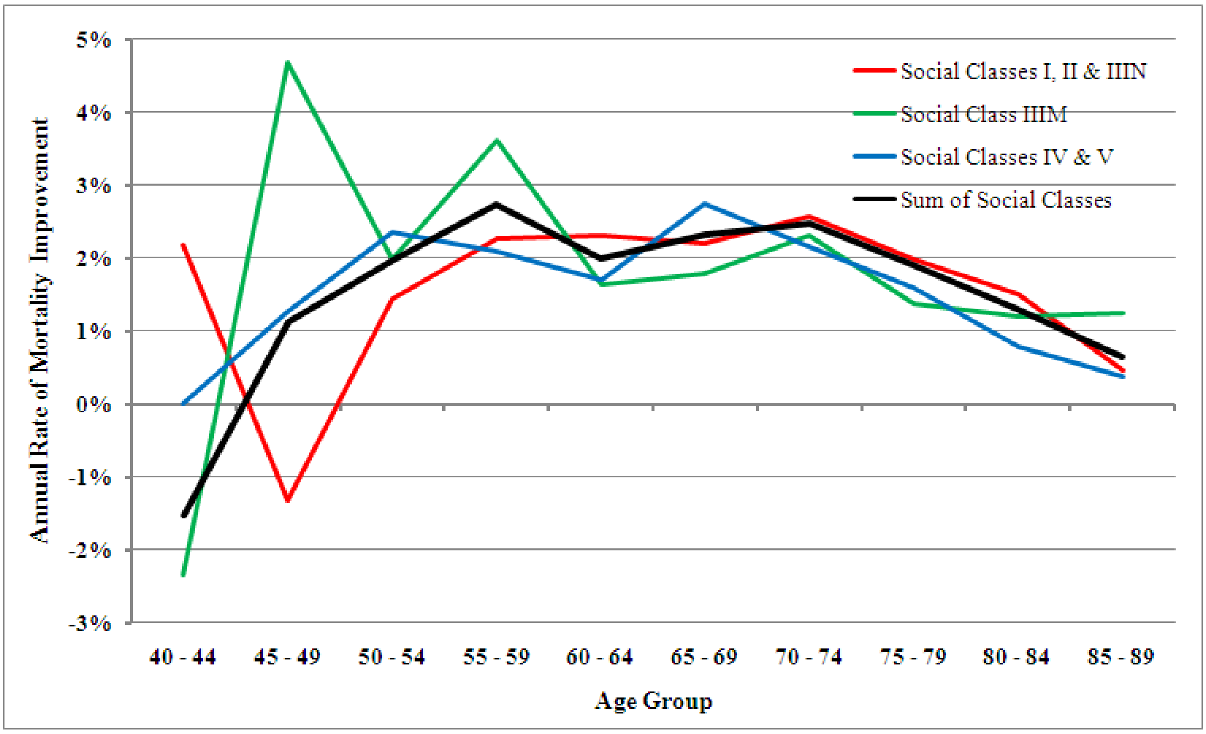

Figure 3.26: Average annual rate of mortality improvement for males in England & Wales, 1987-91 to 2002-05, by age group and social class

In summary, whilst there is longer-term evidence of widening mortality differentials by socio-economic group, the more recent picture is unclear.

3.4. Proposals for the Core Parameter Layer of the CMI Mortality Projections Model

3.4.1. Initial Rates of Mortality Improvement

The Model structure allows Initial Rates of Mortality Improvement to be set by individual age and year-of-birth for Age/Period and Cohort Components respectively. In addition, further years of past mortality improvement rates may be entered to bridge any gap between the Base Mortality Rates table used and the time point for the Initial Rates.

Users of the Advanced parameter layer may set the Initial Rates of Mortality Improvement at this granular level to reflect whatever level and pattern of rates they view as appropriate.

In principle, this parameter group is the one best supported by research and analysis, and is likely to be the one where the flexibility of the Model's parameter sets is used most.

Whilst Users are encouraged to undertake their own analysis, the Working Party consider that the research it has undertaken to date does give a reasonable basis on which to establish a set of default values for the Core parameter layer.

The proposed tables for Initial Rates of Mortality Improvement cover calendar years 1991 to 2005 and contain values for individual ages, separately for males and females. These rates represent the total rate of improvement by age, year and gender before any split into Age/Period and Cohort Components and are referred to in this paper as 'aggregate' rates.

These rates of improvement were derived using an age-cohort P-Spline model fitted to ONS data for the population of England & Wales, for ages from 18 to 102, for the period 1961 to 2007. As well as providing the smoothed data from which to estimate 'current' rates, this approach automatically also provides rates for earlier years on a consistent basis.

The 'current' (initial) rates of mortality improvement are taken as those for calendar year 2005; the first year of the projection is therefore assumed to be 2006. This reflects the Working Party's view that 2005, i.e. 2 years inside the edge of the available data, is the latest year for which sufficiently robust estimates of rates of mortality improvement may be made at present.

Population data has been used for the Core parameter level, primarily because the Working Party considers the estimation errors to be lower than for other datasets. In addition, there did not seem to be a single alternative dataset that would meet the needs of all users. Whilst it is clear that only the population-level data could support an assumption set with strong features by age, it was initially expected that comparison of these smaller (but arguably more relevant) datasets against the population dataset might provide evidence for a high-level modification of the population-based table of rates of mortality improvement. However, the results of this analysis were inconclusive particularly for the most recent decade or so.

An age-cohort P-Spline model with 4-year knot spacing was selected for deriving the default rates. For both males and females, this version of the P-Spline family proved to be the best model based on statistical criteria. More generally, basing its conclusions on the work carried out comparing different P-Spline models and a number of alternative methodologies for smoothing the data, the Working Party consider that this model provides a better representation of the features of the data. In addition, tests were carried out to check, for example, that the resulting projections are not overly sensitive to the choice of smoothing model or methodology.

The proposed Initial Rates of Mortality Improvement for calendar year 2005 are shown in Figure 3.27.

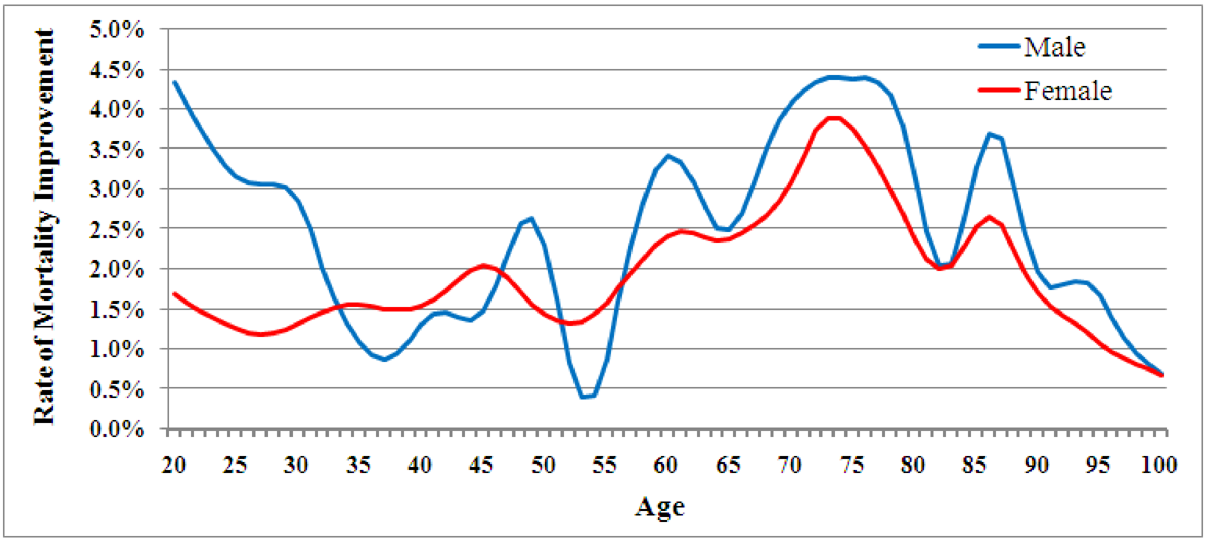

Figure 3.27. Estimated aggregate Initial Rates of Mortality Improvement; 2005

Given the nature and definition of the underlying population data, the rates of mortality improvement derived apply from one calendar year to the next. So, the Initial Rates reflect the falls in mortality rates from $q_{x,2004}$ to $q_{x,2005}$ for lives aged x exact on 01/01/2004 and 01/01/2005 respectively.

The P-Spline model was fitted to data for ages 18 to 102, but the derived rates of improvement were only used for ages 20 to 100 to reduce the possible distortion of 'edge effects' by age.

A very simple run-off was applied above age 100: the aggregate rate of mortality improvement was extrapolated by age (within each calendar year) towards zero at the rate of 0.1% p.a. each year. For example, if the rate at age 100 were 0.5%, the extrapolation would give 0.4% at age 101, 0.3% at age 102, and so on until 0.0% at ages 105 and above.

3.4.2. Split of Initial Rates into Age/Period Component and Cohort Component

The Model design incorporates splitting rates of mortality improvement into Age/Period and Cohort Components. This concept is well-supported by the Working Party's research and has some simple parallels with the 'base projection + cohort' structure of the CMI Interim Cohort Projections, and also with the approach taken by the GAD / ONS for the mortality element of the most recent National Population Projections. This section describes the method the Working Party developed to estimate the split of aggregate Initial Rates of Mortality Improvement into the two components.

Although the task is far from straightforward, the Working Party consider that estimating this split – so that the rate for an individual age might be allocated say 40% to Cohort and 60% to Age/Period Components – does lead to a simpler model in concept, as it removes the ‘missing triangles’ and the old-age complications which occur in models which assume or permit only a sudden switch from by-age to by-cohort projection – that is models in which the rate for each age is allocated either 0% or 100% to the Cohort Component.

A simple additive age-period-cohort (APC) model was created specifically for this purpose. Its inputs were the aggregate Initial Rates of Mortality Improvement, i.e. smoothed aggregate rates of mortality improvement taken from the age-cohort P-Spline model fitted to the ONS dataset for the population of England & Wales.

Two constraints were applied to the APC model to ensure a unique solution. These were:

- The sum over all ages of the age components must be zero; and
- The sum over the full age-year data grid of the cohort components must be zero.

Basis-splines with 4-year knot spacing (consistent with the knot-spacing in the P-Spline models applied to the raw population data) were used to ensure reasonable smoothness in the derived APC components (but without imposing any particular shape).

The APC model was fitted to the smoothed aggregate rates to estimate the age, period and cohort components separately (utilising ‘solver’ in Excel to minimise the sum of squares of residual errors, having checked the solution produced was unique and stable).

The age and period elements could be combined to give the Age/Period Component, and the cohort element mapped directly to the Cohort Component. However, it is also necessary to assess and allocate the residual error, i.e. the difference between the aggregate rates from the P-Spline model and the sum of the APC model elements.

Figure 3.28 shows the period components derived from the APC model and compares them against the actual (all age) average rate of improvement by calendar year. The close match between the APC model period component and the actual averages reflects the constraints applied.

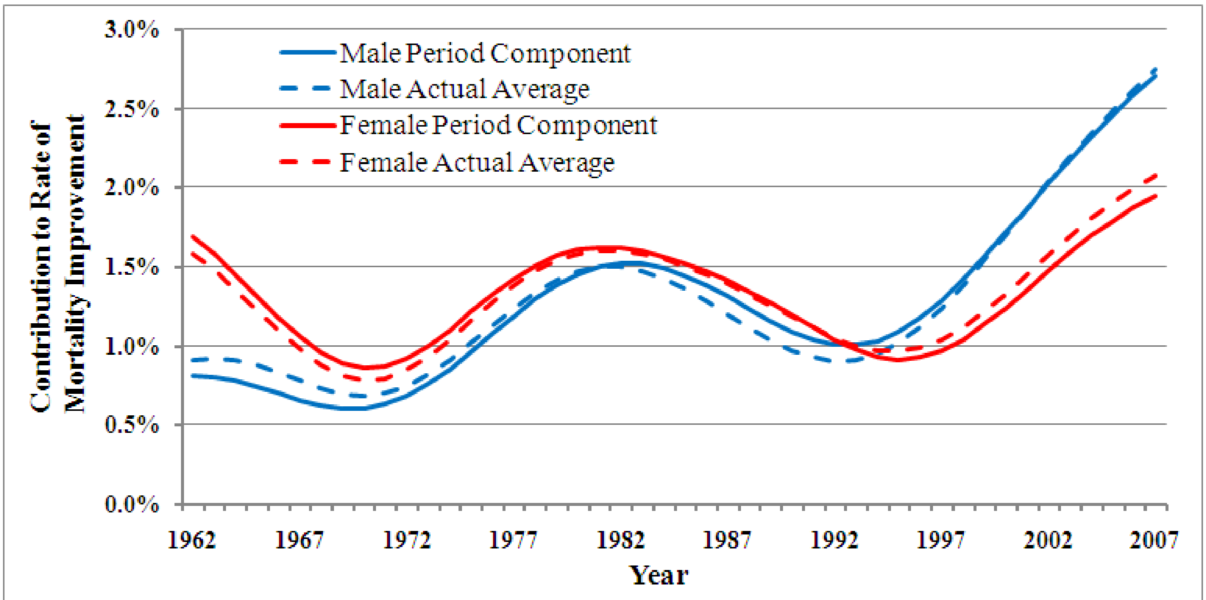

Figure 3.28: Comparison of Period Components derived in APC Model against All Age Average Rates of Improvement from Input Table

Figures 3.29 and 3.30 show the age, period and cohort components from the fitted APC models, for males and females respectively. Also shown are the residual errors.

Reasonableness checks were performed on the age, period and cohort components, for example by comparison against the Working Party's previous analysis of the underlying population data. The Working Party consider that the peaks and troughs of the derived cohort component reasonably reflect those noted in earlier analysis, and that the raised age component for age 60 to 80 reasonably reflects the feature of generally high rates of mortality improvement for those ages over the past two decades (most likely driven by reductions in mortality rates from circulatory disease).

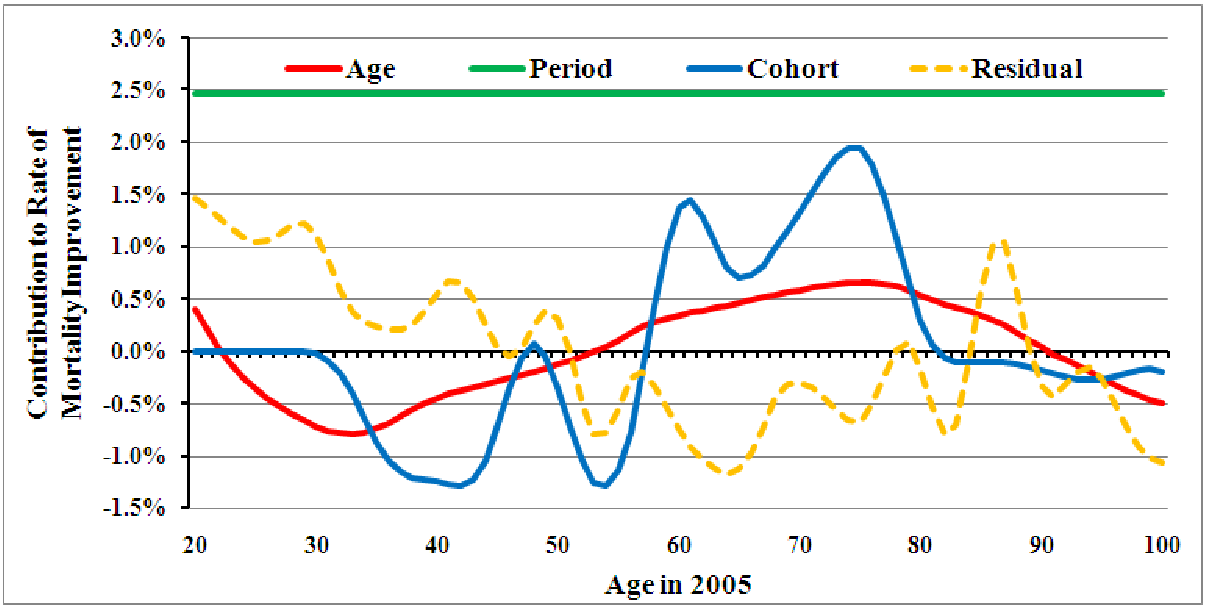

Figure 3.29: Results Derived from Fitted APC Model, Males; 2005 Age, Period and Cohort Components, plus Residual Errors

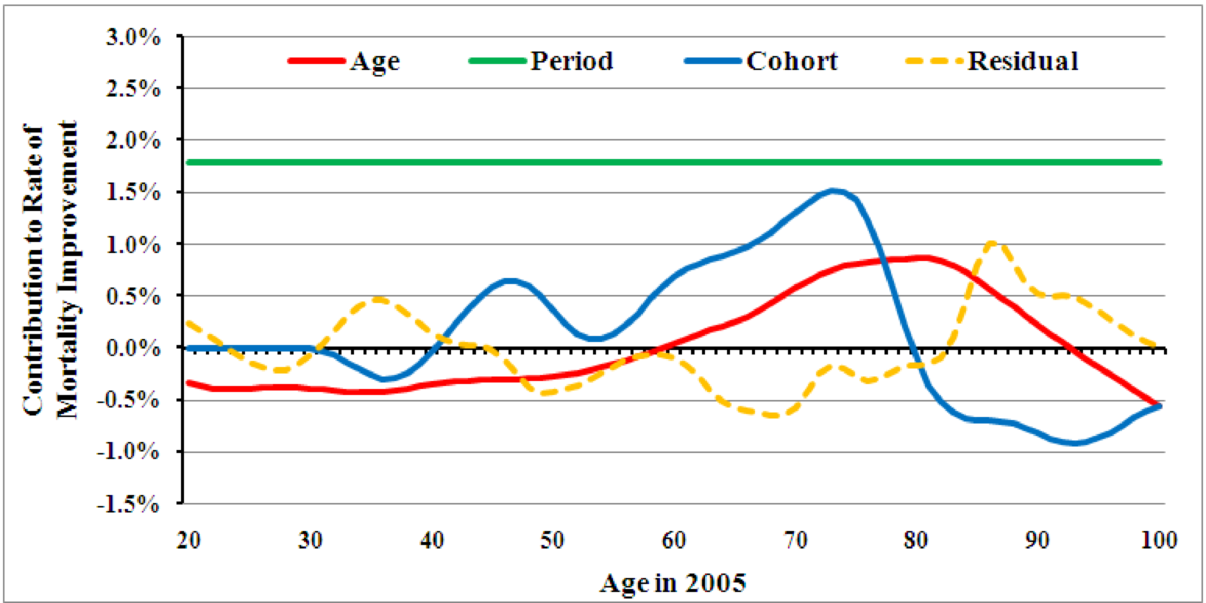

Figure 3.30: Results Derived from Fitted APC Model, Females; 2005 Age, Period and Cohort Components, plus Residual Errors

The Working Party initially considered allocating all the residual errors to the Age/Period Component on the grounds that these were likely to be short-term deviations from the underlying model and so should be run-off relatively quickly. However, this was subsequently amended at the middle and higher ages for the following reasons:

- At the higher ages, the default Period of Convergence parameters in the Model run-off the Age/Period Component more slowly than the Cohort Component.
- Residual differences seen in the fit against the underlying P-Spline Model appeared as short-term cohorts, making the Age/Period allocation look somewhat inconsistent.
- Allocation to the Age/Period Component produced a visible spike at certain high ages. This was not considered to be a genuine feature of the data and so was deemed inappropriate.
- It is generally accepted that for middle and higher ages, the mix of causes of deaths changes slowly with age. This led to the conclusion that this spike, and any other narrow-age-range fluctuations, should be assigned to the Cohort Component rather than the Age/Period Component for middle and higher ages.

The Working Party decided to allocate the residual errors below age 30 to the Age/Period Component, and errors above age 60 to the Cohort Component, with a linear transition in between. In addition, the fitted Cohort Component was constrained to be zero near the edges of the data and in particular up to age 30 (as there are too few years' data, and too much 'noise' in the data at young ages, to form a safe conclusion on cohort components).

Figures 3.31 and 3.32 show the final derived Age/Period and Cohort Components of the Initial Rates of Mortality Improvement.

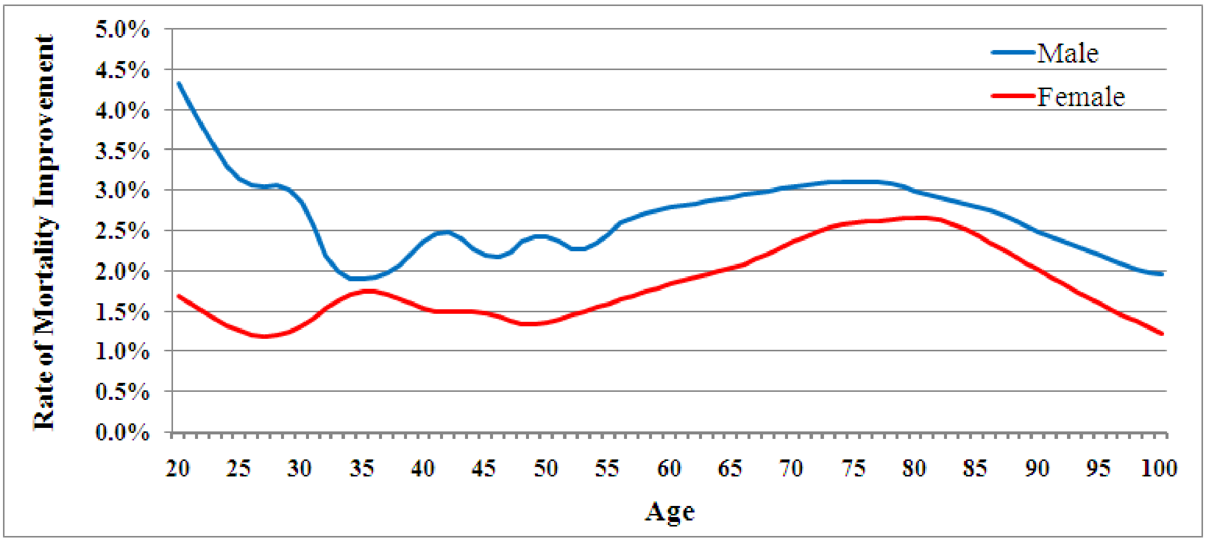

Figure 3.31 Estimated Age/Period Component of Initial Rates of Improvement; 2005

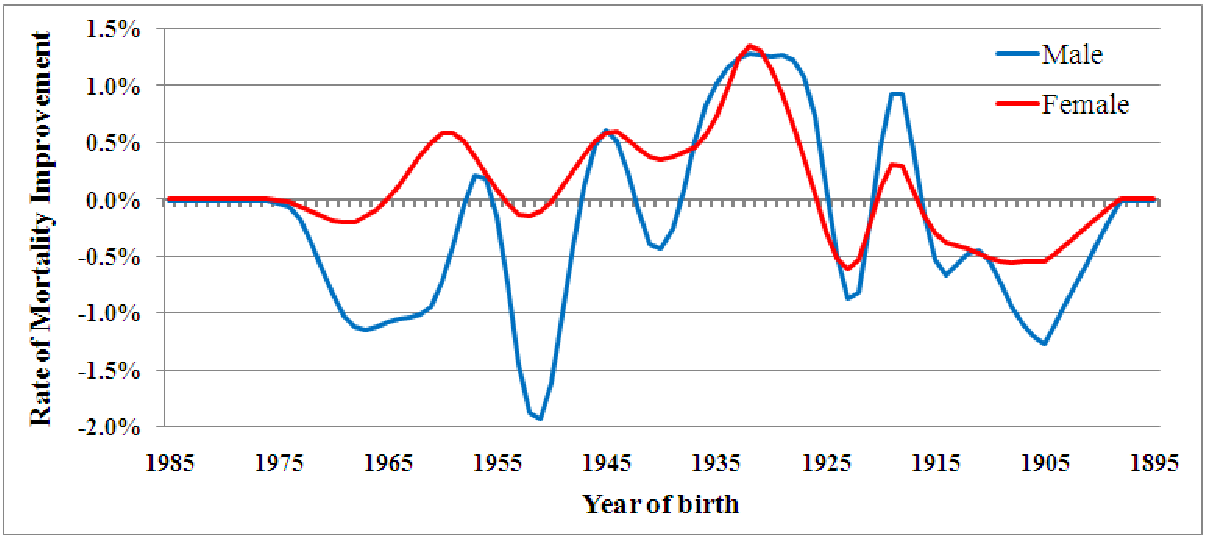

Figure 3.32: Estimated Cohort Component of Initial Rates of Improvement; 2005

The Working Party noted the negative strength of the cohort component for males between ages 30 and 40 (e.g. -1.2% at age 38) but satisfied itself through a review of the evidence, including heat maps, that this is a genuine cohort feature and consider it not unreasonable to project it on a cohort basis.

4. Long-Term Rates of Mortality Improvement

4.1. Long-Term Rates of Mortality Improvement: Past Experience

4.1.1. High-Level Trends in Mortality Rates: Seven developed countries; 1854-2004

A view on suitable Long-Term Rates of Mortality Improvement may be usefully informed by an analysis of past trends in mortality improvement over a long period of time. The Working Party's analysis of an international average rate of improvement was performed using experience data from seven countries included in the Human Mortality Database (www.mortality.org) in December 2008. The seven countries (Belgium, Denmark, England & Wales, France, Netherlands, Norway and Sweden) were selected as they each offered time-series data for a period of at least 150 years.

Figure 4.1 shows how average annual rates of improvement for ten-year age bands for males changed over successive periods of 25-years. To calculate the rate of improvement for a given age, the average mortality rate was calculated for the three-year period centred on the beginning and end of each period. For instance the improvement for the period 1979 to 2004 is the average annual rate of change based on a mortality rate calculated for the period 2003-2005 compared with one calculated for the period 1978-1980.

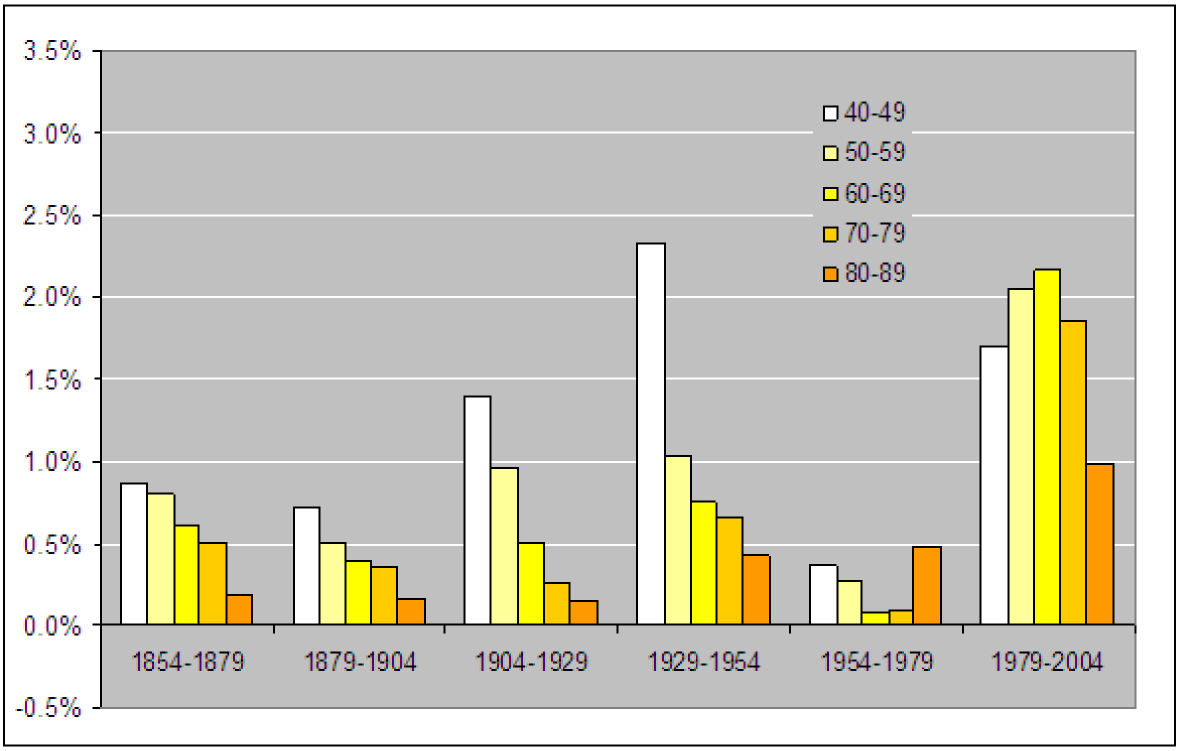

Figure 4.1: Average annual rate of improvement for males in Belgium, Denmark, England & Wales, France, Netherlands, Norway and Sweden, successive periods of 25 years, by age group

Whilst mortality improvement rates averaged over very long periods tend to even out across ages and gender, the patterns of mortality change do shift materially even when comparing quarter-century intervals.

It is particularly notable that rates of improvement for males have been far higher in the last 25 years than in any other period for most ages. The average rate of change (for ages 40 to 89) has been 1.8% p.a. for the most recent 25-year period. The equivalent average for the previous 125 years was 0.6% p.a.

It is also apparent that the rapid improvements experienced in the last 25 years followed a quarter-century over which there was little change in mortality rates for ages 40 to 89 (i.e. an average of 0.3% p.a.).

The change in mortality improvement rates for males from 1954-1979 to 1979-2004 is likely to reflect a substantial 'drag down' from the effects of smoking in the earlier period, and then a corresponding increase in the rate of mortality improvements over the last 25 years as the falling prevalence of smoking has fed through, along with general falls in cardiovascular mortality.

To those familiar with the current rates of change in UK mortality, even the averages for the last 25 years appear low. The average default Initial Rate of Improvement for males aged 40-89 in the prototype Model, derived using a P-spline model fitted to England & Wales population data, is 2.7% p.a.

Figure 4.2 presents the results of the analysis in a slightly different way; with average rates of change over successive periods of 50 years plotted for each individual age between 20 and 100.

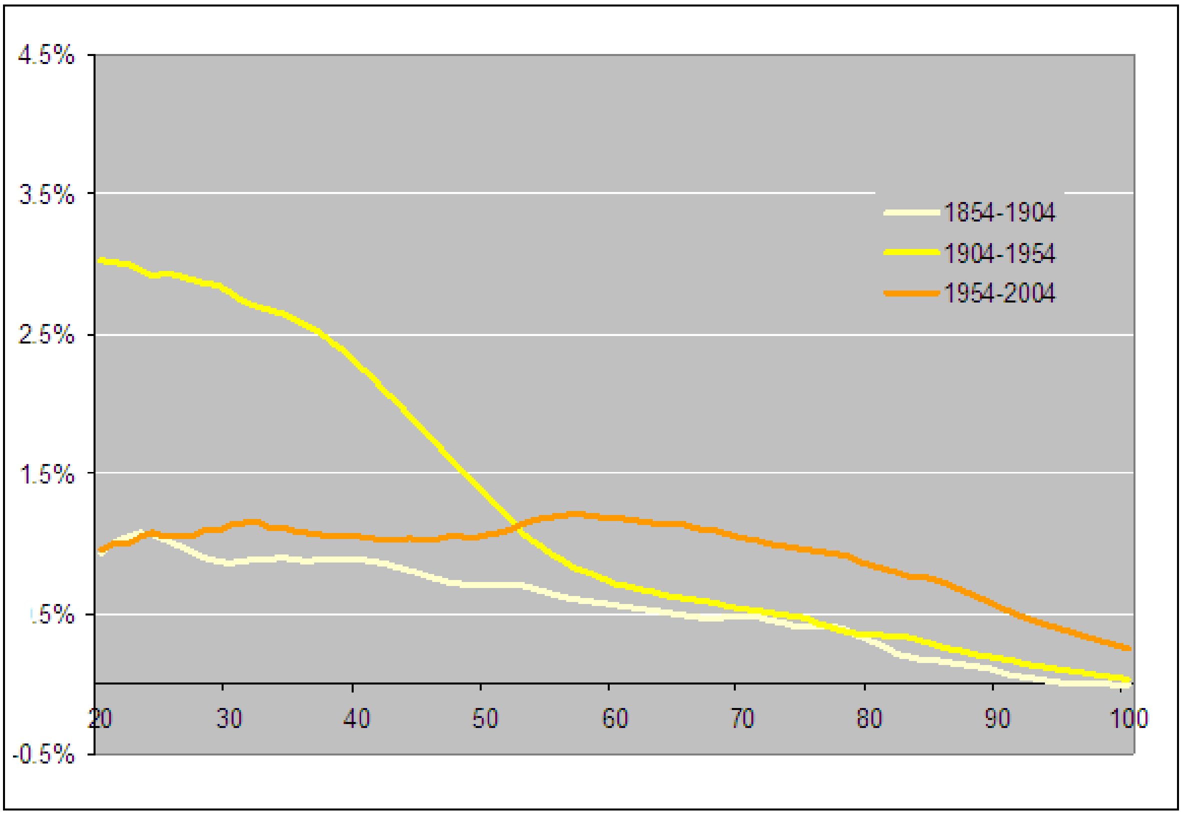

Figure 4.2: Average annual rate of improvement for males in Belgium, Denmark, England & Wales, France, Netherlands, Norway and Sweden, successive periods of 50 years, by age

It is clear from figure 4.2 that, other than over the last 50 years, improvements for males have mostly been concentrated at younger ages, with very little change above age 60.

Over the second half of the 19th century, average mortality improvement rates fell almost linearly with age over the range 20-100, but were relatively modest at all ages.

Average rates of improvement over the next 50 years were far higher at younger ages, particularly for men aged in their 20s and 30s, but were similar to rates for the previous half century at older ages.

Over the last 50 years, average rates of improvement have been similar for all ages under 70 (averaging 1.1% p.a.), but have reduced with increasing age.

Figure 4.3 compares average rates of change, by age, in the seven countries referred to above with an average from a larger group of nineteen countries, for whom data was available on the Human Mortality Database for a period of at least 50 years.

This comparison suggests that the improvements experienced by the group of seven countries are not untypical of a broader set of developed nations, at least over the last 50 years.

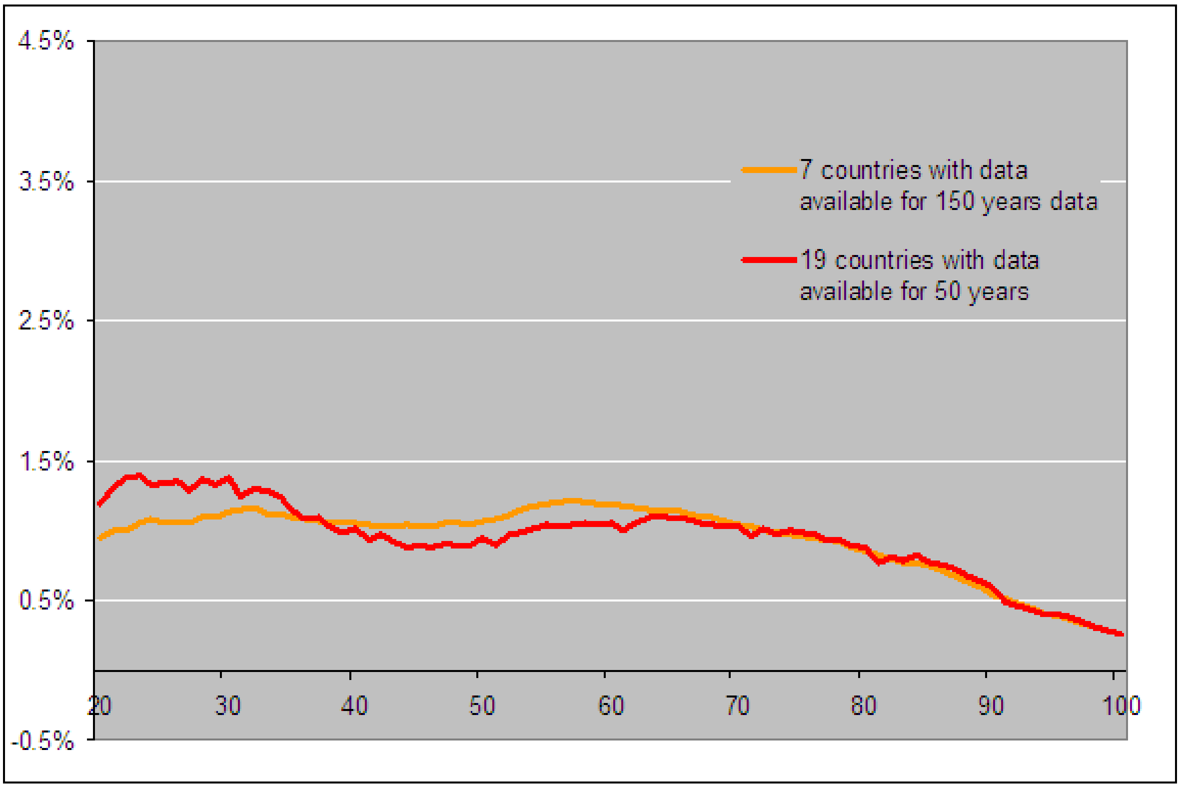

Figure 4.3. Average annual rate of improvement for males over successive periods of 50 years, by age, various groups of countries

Figure 4.4 shows the equivalent average rates of improvement for females for the same group of seven countries.

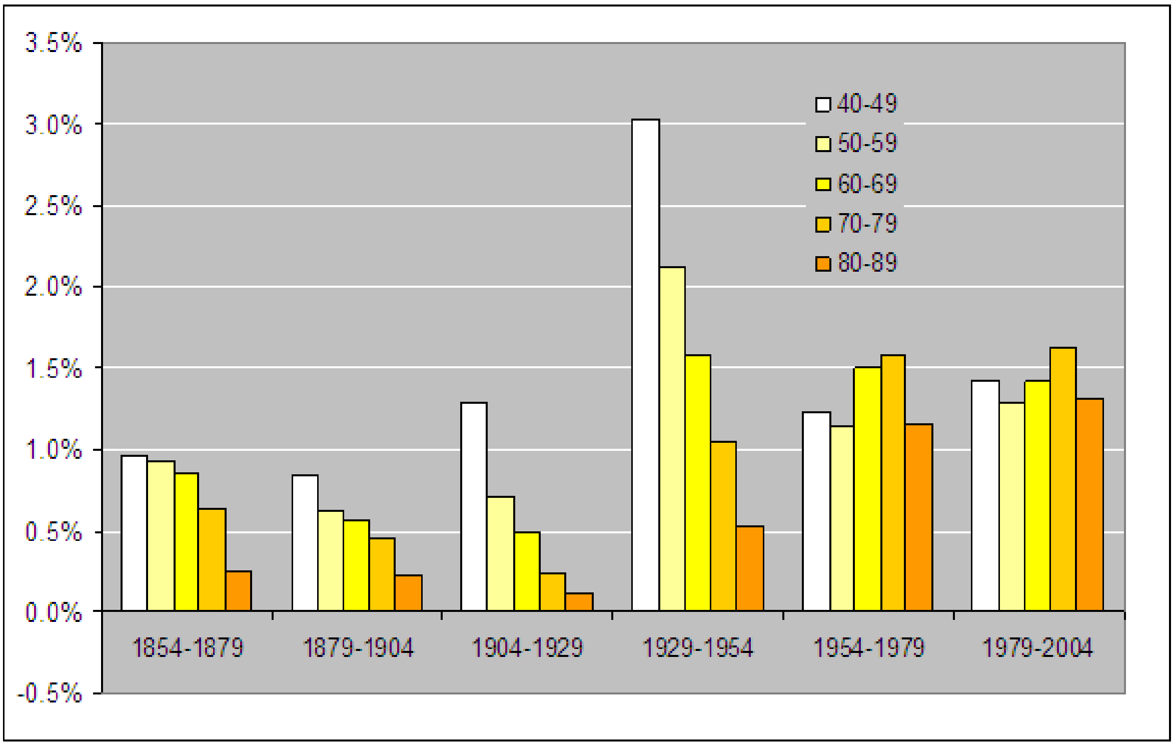

Figure 4.4. Average annual rate of improvement for males in Belgium, Denmark, England & Wales, France, Netherlands, Norway and Sweden, successive periods of 25 years, by age group

Figure 4.4 shows that average improvements for females have been of a similar magnitude over the last 25-year period compared with those over the preceding 25 years. In fact the rate of average rate of change for females aged 40-89 was 1.7% over 1929-1954; 1.3% over 1954 to 1979 and 1.4% over 1979 to 2004. This relatively even distribution over time contrasts strongly with the equivalent Figures for males. This reflects the fact that changes in smoking prevalence and reductions in cardiovascular mortality have had less influence on trends in female mortality compared to male mortality rates.

Figure 4.5 shows a similar picture to Figure 4.2, for males. The pace of improvement has generally reduced with increased age. However, the rates of improvement for younger females were particularly rapid during the period from 1904-1954 and the improvements for older females have been significantly more rapid in the last 50 year-period than in preceding periods.

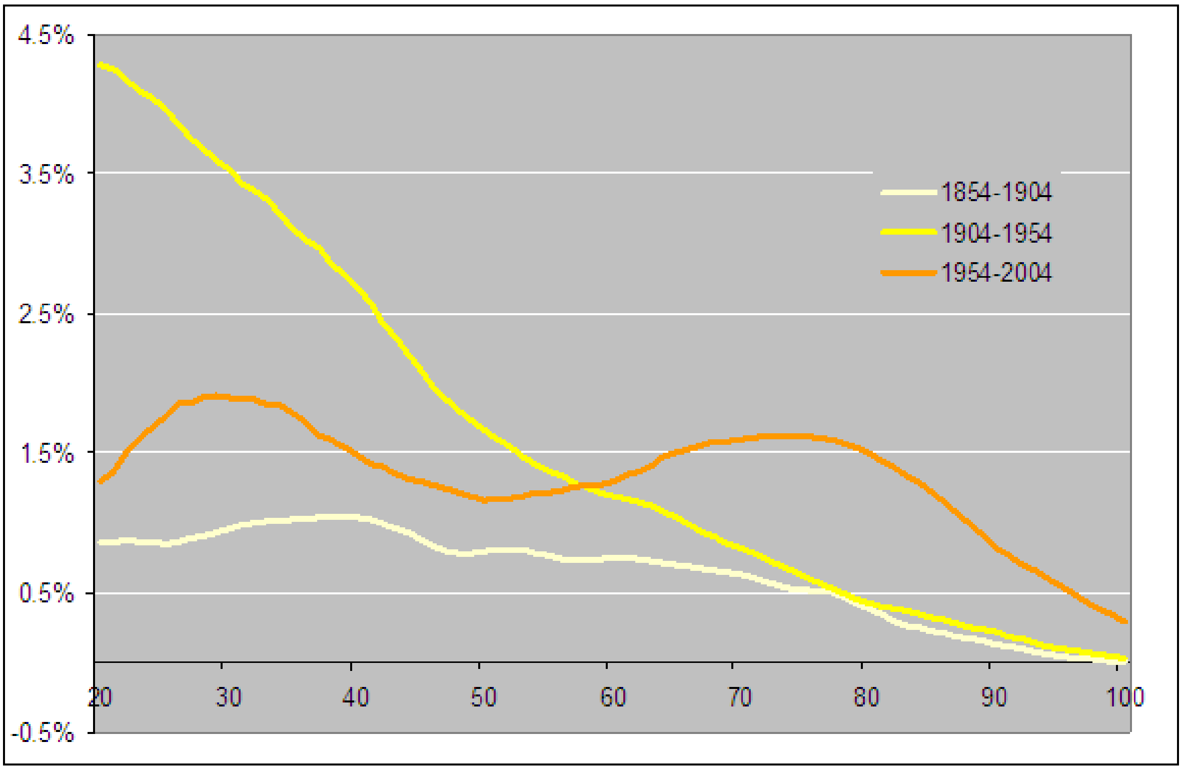

Figure 4.5. Average annual rate of improvement for females in Belgium, Denmark, England & Wales, France, Netherlands, Norway and Sweden, successive periods of 50 years, by age

Figure 4.6 compares average rates of change for females, by age, in the seven countries referred to above with the average from the larger group of nineteen countries.

As was the case for males, this comparison suggests that the improvements experienced by the group of seven countries are not untypical of a broader set of developed nations over this period.

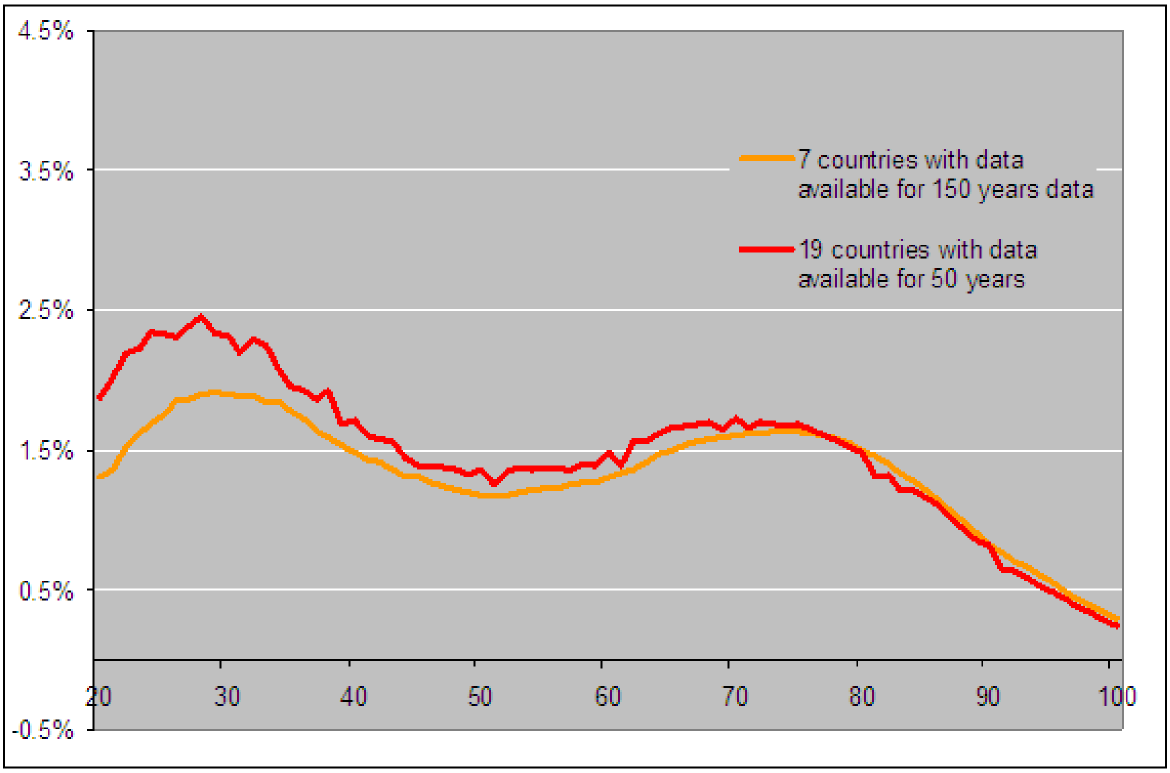

Figure 4.6. Average annual rate of improvement for females over successive periods of 50 years, by age, various groups of countries

4.1.2. High-Level Trends in Mortality Rates: England & Wales Population; 1854-2004

Figures 4.7 and 4.8 show values for England & Wales equivalent to those given in Figures 4.4 and 4.5 for the average of seven countries (Belgium, Denmark, England & Wales, France, Netherlands, Norway and Sweden).

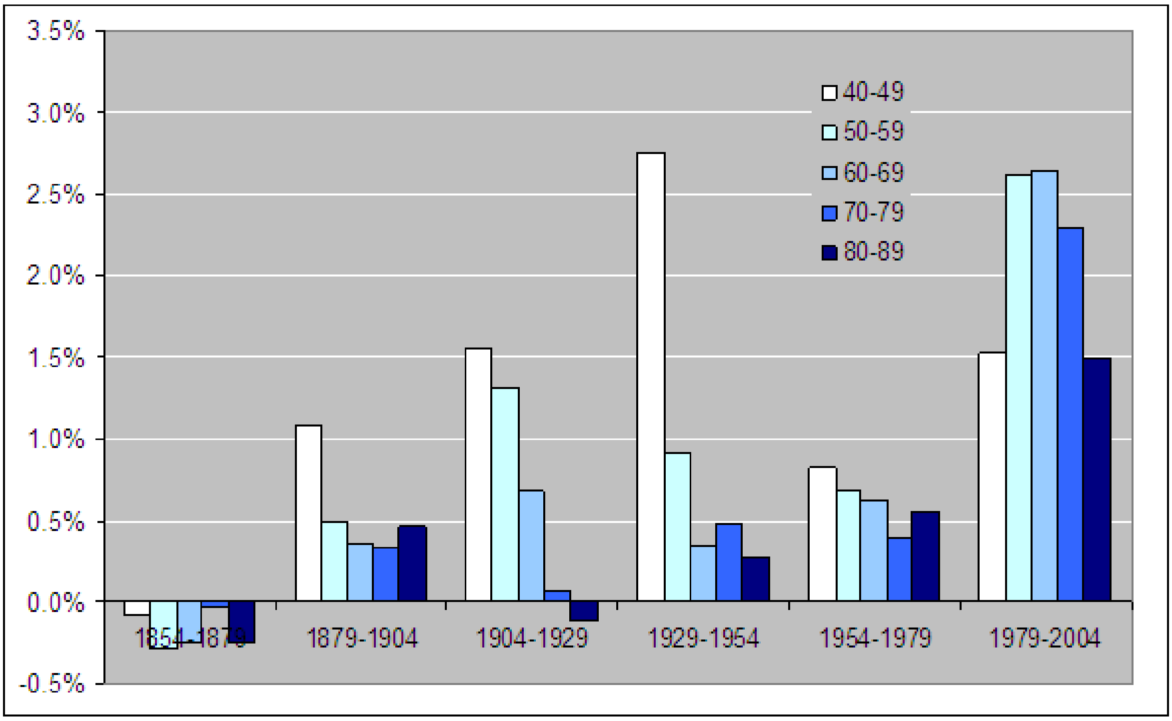

Figure 4.7. Average annual rate of improvement for males in England & Wales, successive periods of 25 years, by age group

The values in Figure 4.7 indicate that the rates of change for males aged 40-89 over the last 25 years have been somewhat higher in England & Wales than the international average (2.1% p.a. versus 1.8% p.a.). The equivalent Figure for the preceding 125 years shows a slightly lower average in England & Wales (0.5% p.a. versus 0.6% p.a.).

Over the entire 150-year period the average rates of change have been very similar, as illustrated by Table 4.1.

Table 4.1. Average annual rate of improvement for males in England & Wales and an average of seven countries, 1854-2004, by age group

Age group	Average annual rate of mortality improvement	
	England & Wales	International average (seven countries)
40-49	1.3%	1.2%
50-59	1.0%	0.9%
60-69	0.7%	0.8%
70-79	0.6%	0.6%
80-89	0.4%	0.4%
All ages (40-89)	**0.8%**	**0.8%**

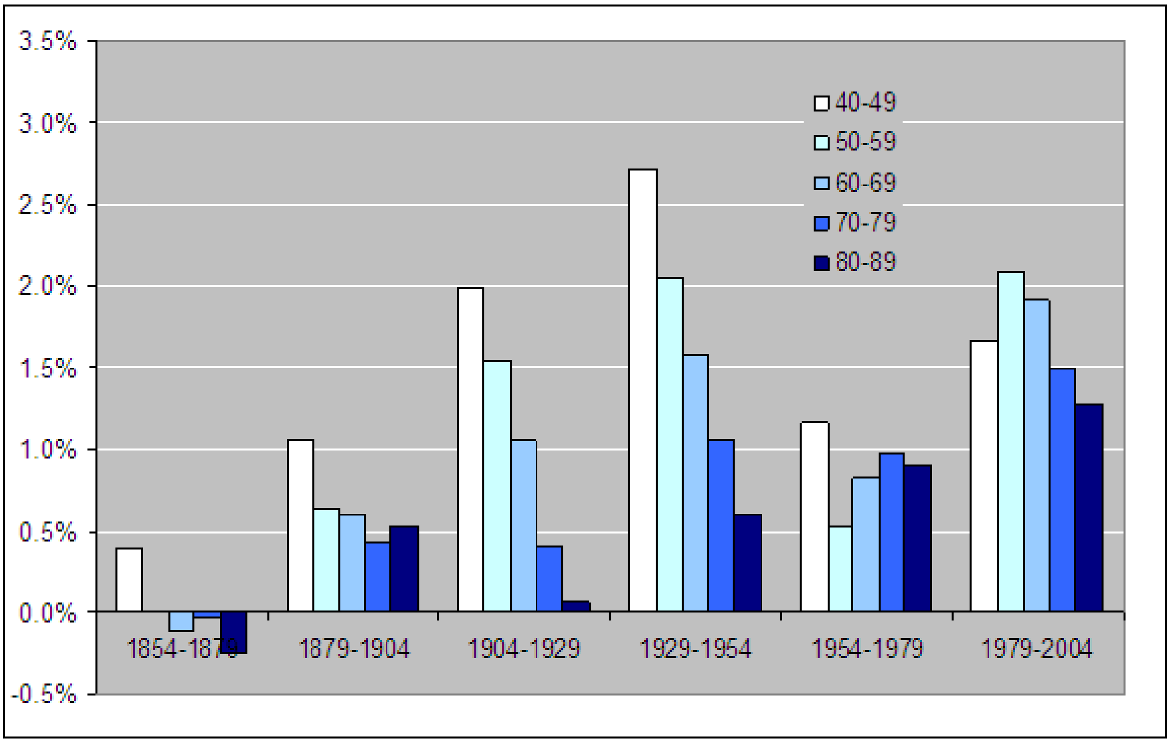

Figure 4.8. Average annual rate of improvement for females in England & Wales, successive periods of 25 years, by age group

The values in shown in Figure 4.8 also indicate that rates of change for females aged 40-89 over the last 25 years have been somewhat higher in England & Wales than the international average (1.7% p.a. versus 1.4% p.a.). The equivalent Figure for the preceding 125 years shows a slightly lower average in England & Wales (0.8% p.a. versus 1.0% p.a.).

Over the entire 150-year period the average rates of change have again been very similar, as illustrated by Table 4.2.

Table 4.2. Average annual rate of improvement for females in England & Wales and an average of seven countries, 1854-2004, by age group

Age group	Average annual rate of mortality improvement	
	England & Wales	International average (seven countries)
40-49	1.5%	1.5%
50-59	1.1%	1.1%
60-69	1.0%	1.1%
70-79	0.7%	0.9%
80-89	0.5%	0.6%
All ages (40-89)	**1.0%**	**1.0%**

4.2. Long-Term Rates of Mortality Improvement: Insured Lives & Pensioners

This section describes an additional analysis performed by the Working Party to investigate whether patterns of change for different sub-groups of the population might justify differential assumptions on long-term rates of mortality improvement.

Average annual mortality improvement rates from the ONS England & Wales population dataset (males and females, 1961-2007) were compared against those from two CMI datasets:

- CMI Permanent Assurances Assured Lives (males 1947-2006, females 1983-2006)
- CMI Pensioners (males and females 1983-2006)

The approach taken for the comparison was to smooth the raw data by fitting P-Spline mortality surfaces to each dataset and using them to derive estimates of annual rates of mortality improvement for individual ages and years. Average annualised improvement rates were then calculated, separately for each age, over various time periods and the results compared between datasets. For all comparisons, care was taken to ensure the same type of P-Spline model and the same knot-spacing were adopted for both datasets.

Figure 4.9 shows the average rates of mortality improvement, over the 40 year period 1965-2004, for the male population and CMI Permanent Assurances datasets (the only pair with sufficient data history).

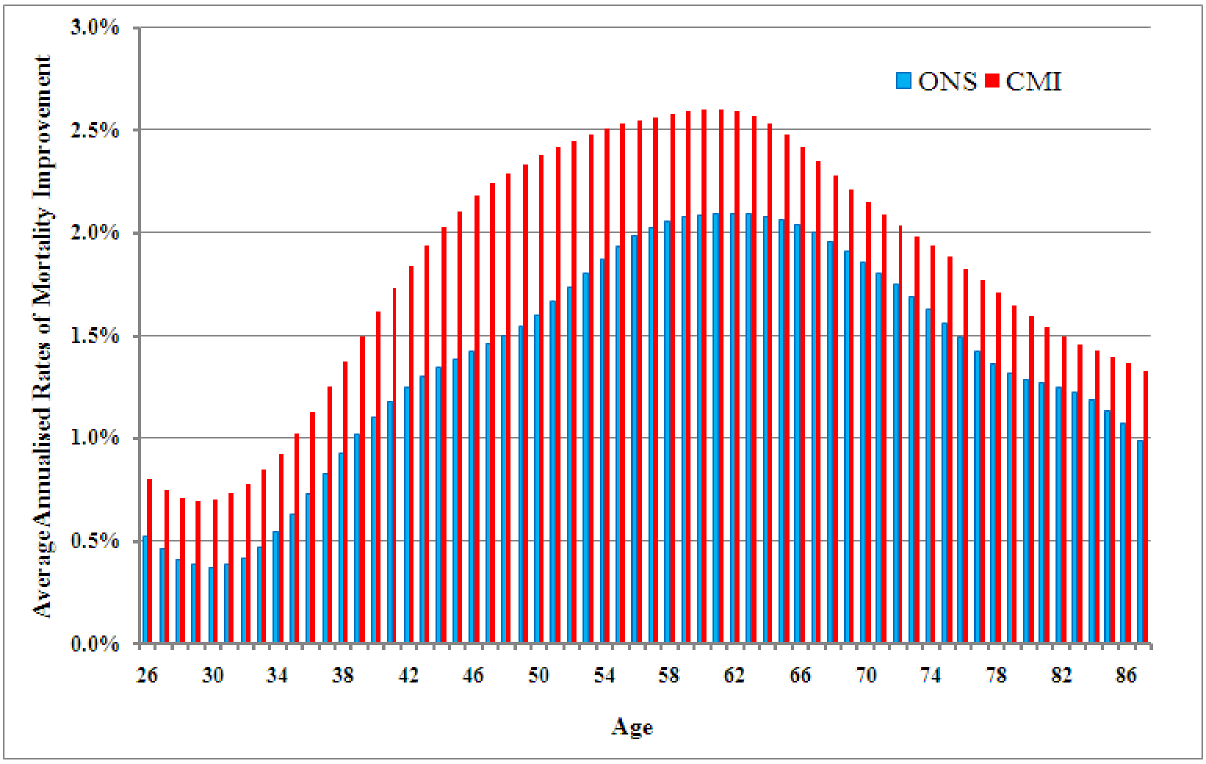

Figure 4.9: Average annualised rates of mortality improvement, 1965-2004
ONS England & Wales population and CMI Permanent Assurances datasets; males

Over the 40-year period as a whole, the CMI Permanent Assurances dataset shows higher rates of mortality improvement than the population data for male lives at all ages. This is in line with the general observation of widening socio-economic differentials over the period as a whole. The pattern by age of the differential in average rates of improvement appears closer to a flat addition to the population rate, rather than (say) a multiple of it. This difference is considered more closely in Figure 4.10.

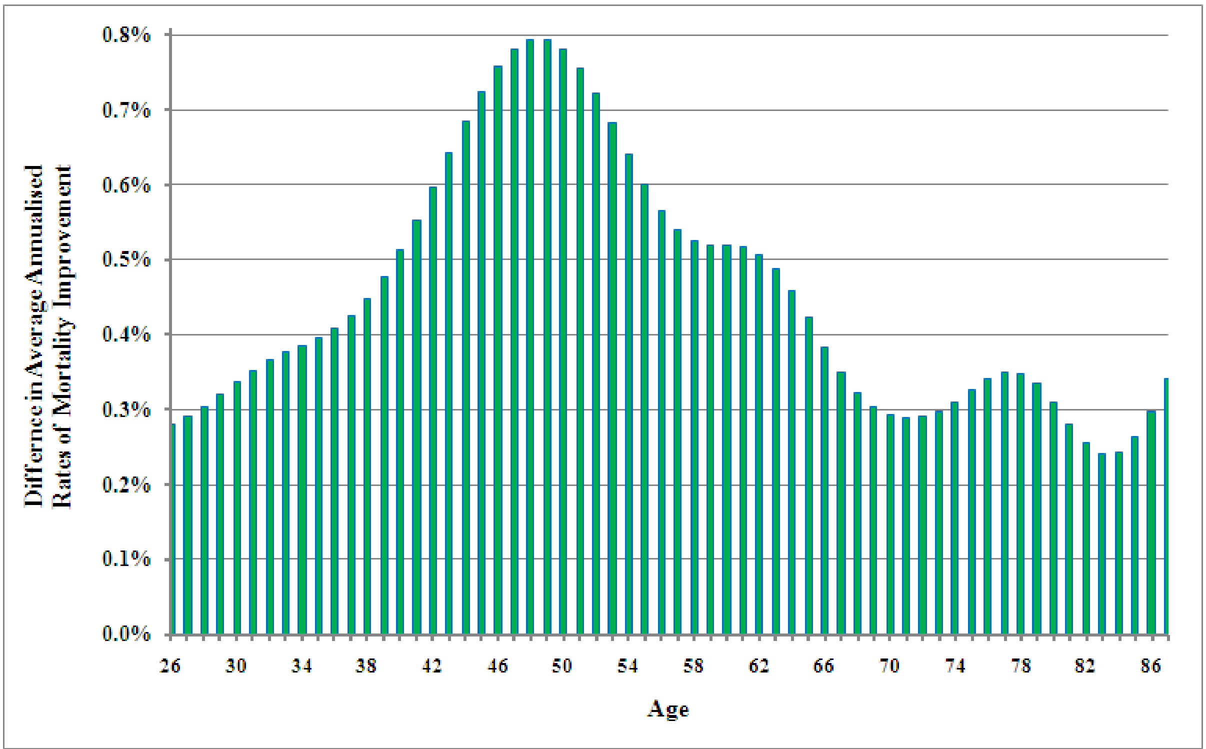

Figure 4.10: Difference in average annualised rates of mortality improvement, 1965-2004, ONS England & Wales population and CMI Permanent Assurances datasets; males (CMI average rate minus ONS average rate for each age)

Figure 4.10 shows that the average extra rate of mortality improvement for the CMI Permanent Assurances dataset over the population data varies between 0.2% and 0.8% p.a. The difference has been greatest (in absolute terms) at ages 35-65, with the peak difference occurring around age 48. As noted previously, however, P-Splines apply a greater degree of smoothing to the lower-volume CMI dataset than to the population data (see section 3.2.2), which may distort the apparent pattern of differentials by age.

Figure 4.11 compares this same pair of datasets but splits the full 40-year period into four 10-year periods, giving a view of the changes over time in the average rates of mortality improvement and of the differential in rates between the datasets.

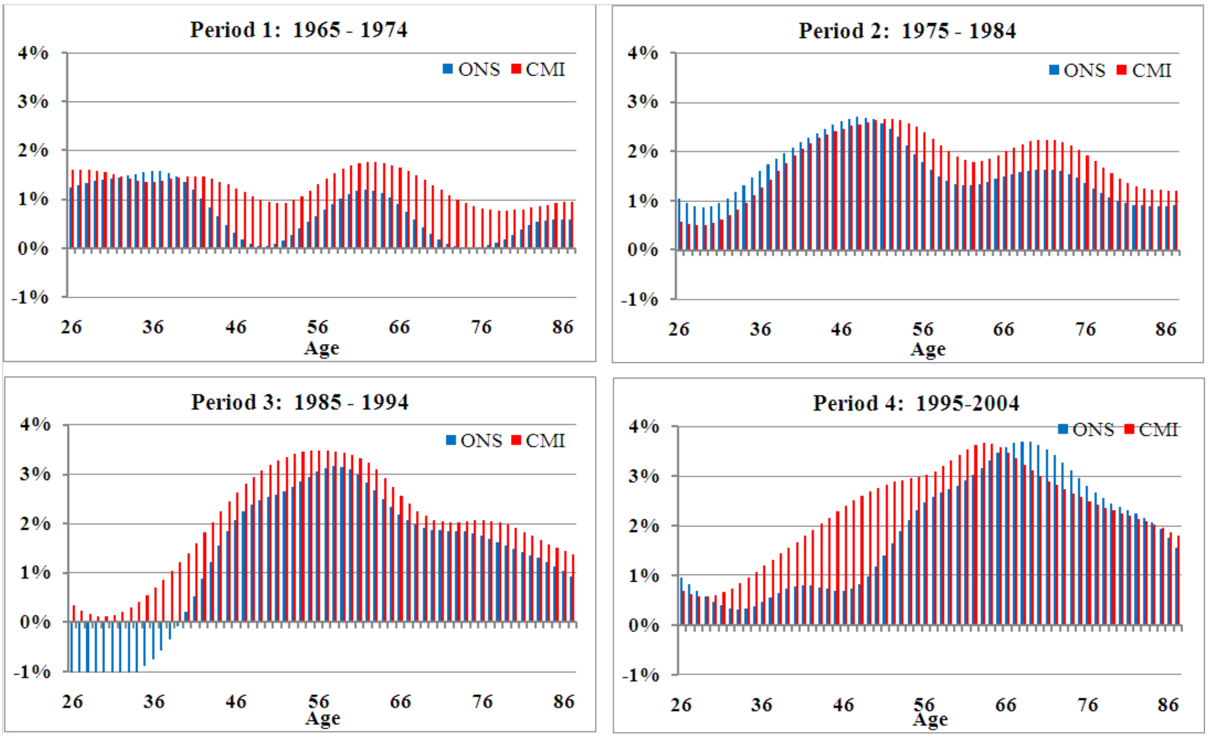

Figure 4.11: Average annualised rates of mortality improvement for four successive 10-year periods covering 1965-2004 ONS England & Wales population and CMI Permanent Assurances datasets; males

Figure 4.11 shows a similar pattern to Figure 4.9 for the first 30 years. Ignoring the youngest ages, where rates are naturally more volatile, we see evidence of a broadly flat differential by age throughout 1965-1994. In the last 10 years, however, the pattern appears to have changed such that mortality has improved more rapidly for the population than for CMI Assured Lives for ages around 65 and above. Again this is in line with trends seen in population data by socio-economic group.

Another cautionary observation, though, is that the CMI Permanent Assurances data volumes have fallen over the last 20 years and have also been subject to changing mix of contributors. These factors may distort the results and certainly reduce confidence in drawing firm conclusions.

Figure 4.12 shows the average rates of mortality improvement, over the 20 year period 1985-2004, for the male population and CMI Life Office Pensioner datasets. The CMI Pensioner data only covers the period from 1983, forcing the restriction to a 20-year analysis period, and analysis has also been limited to ages 60 and above. The available time period has again been split into 10-year periods to look more closely at any patterns over time.

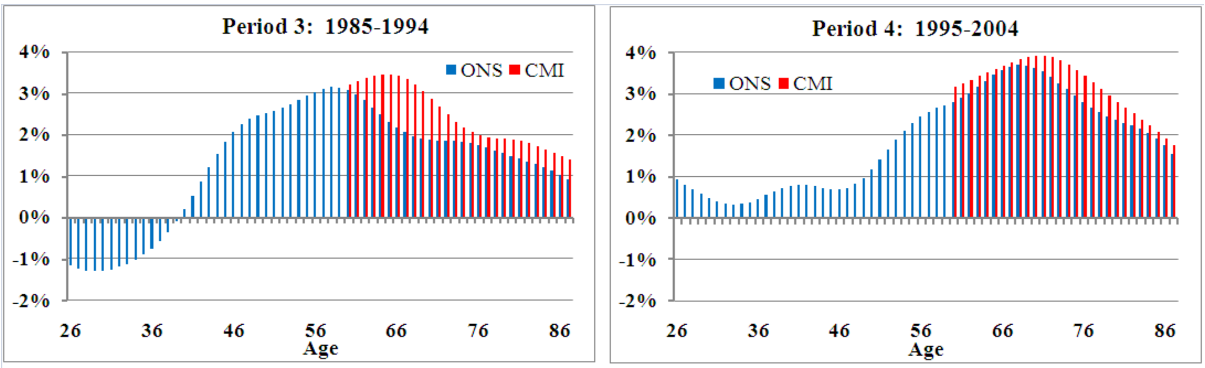

Figure 4.12: Average annualised rates of mortality improvement for two successive 10-year periods covering 1985-2004
ONS England & Wales population and CMI Life Office Pensioner datasets; males

For males aged 60 and above, the CMI Life Office Pensioner dataset shows faster rates of mortality improvement, than the population dataset, throughout 1985-2004. Again, the pattern by age of the differential in average improvement rates appears closer to a flat addition to the population rates rather than a multiple. Whilst the differential may have been reduced in the last 10 years, compared to the previous 10-year period, it has not reversed for the CMI Life Office Pensioners dataset (in contrast to the results for the CMI Permanent Assurances dataset).

To extend the analysis to female lives, Figure 4.13 shows the average rates of mortality improvement, over the 20 year period 1985-2004, for the female population and CMI Permanent Assurances datasets. The available time period has again been split into two 10-year periods to look more closely at any patterns over time.

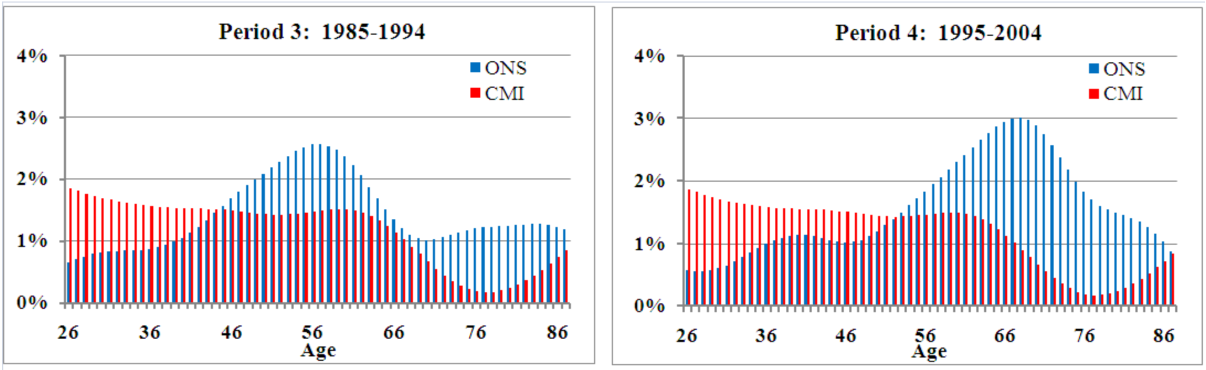

Figure 4.13: Average annualised rates of mortality improvement for two successive 10-year periods covering 1985-2004
ONS England & Wales population and CMI Permanent Assurances datasets; females

This comparison of Insured and population datasets for female lives shows a more erratic picture. However, we note that the average annualised rates of mortality improvement are higher for the population dataset, than the CMI Permanent Assurances dataset, from around age 50. No clear difference can be seen between the most recent and preceding 10-year periods.

4.3. The Range of Long-Term Rates of Improvement Implicit in Published Projections

4.3.1. Mortality Projections in the UK

Figure 4.14 shows the average rates of mortality improvement implicit in a large sample of the projections contained within the CMI Library of Mortality Projections v1.1.

For each of the sample published projections, average annual rates of mortality improvement have been calculated from the projected mortality reduction factors over the period 2030 to 2055. This period was chosen as broadly representative of the expected principal time horizon for the Long-Term Rate parameter in the Model. The calculations were carried out for each individual age, and a further average of the resulting rates, this time across ages, was taken for each of the four 20-year age groups shown in the chart.

The sample projections, chosen as a representative sample of those in the Library, are:

- Medium and Long Interim Cohort Projections, and a further combination of the two (with a 1.5% pa floor)
- The Principal, Higher and Lower variants from the 2006-based National Population Projections for England & Wales, published by ONS
- Lee-Carter models applied to the population dataset for England and Wales (data to 2007) and to the CMI Permanent Assurances dataset (data to 2006)
- Age-period and age-cohort P-Spline models, each applied to the population dataset for England and Wales and to the CMI Permanent Assurances dataset.
    - P-Spline model projections vary far more those of the Lee-Carter model when additional years' data are brought into the model. Results are shown for population dataset up to 2003, 2005 and 2007, and CMI datasets up to 2003 and 2006.

Where possible, results are shown for both males and females. However, the Interim Cohort Projections and projections based on the CMI Permanent Assurances dataset are shown only for males.

The implicit average rates of mortality improvement for each of the various projections are shown by markers of different shapes and colours.

Whilst the results for individual projections may be followed through if desired, the main intention of presenting data in this form is to provide a general picture of the range of implicit long-term rates, and some indication of any variation by age and gender, and by type of projection model. It is clear that long-term rates of improvement implicit in different projections vary enormously; the difference between the Interim Cohort Projections and P-spline projections being particularly marked.

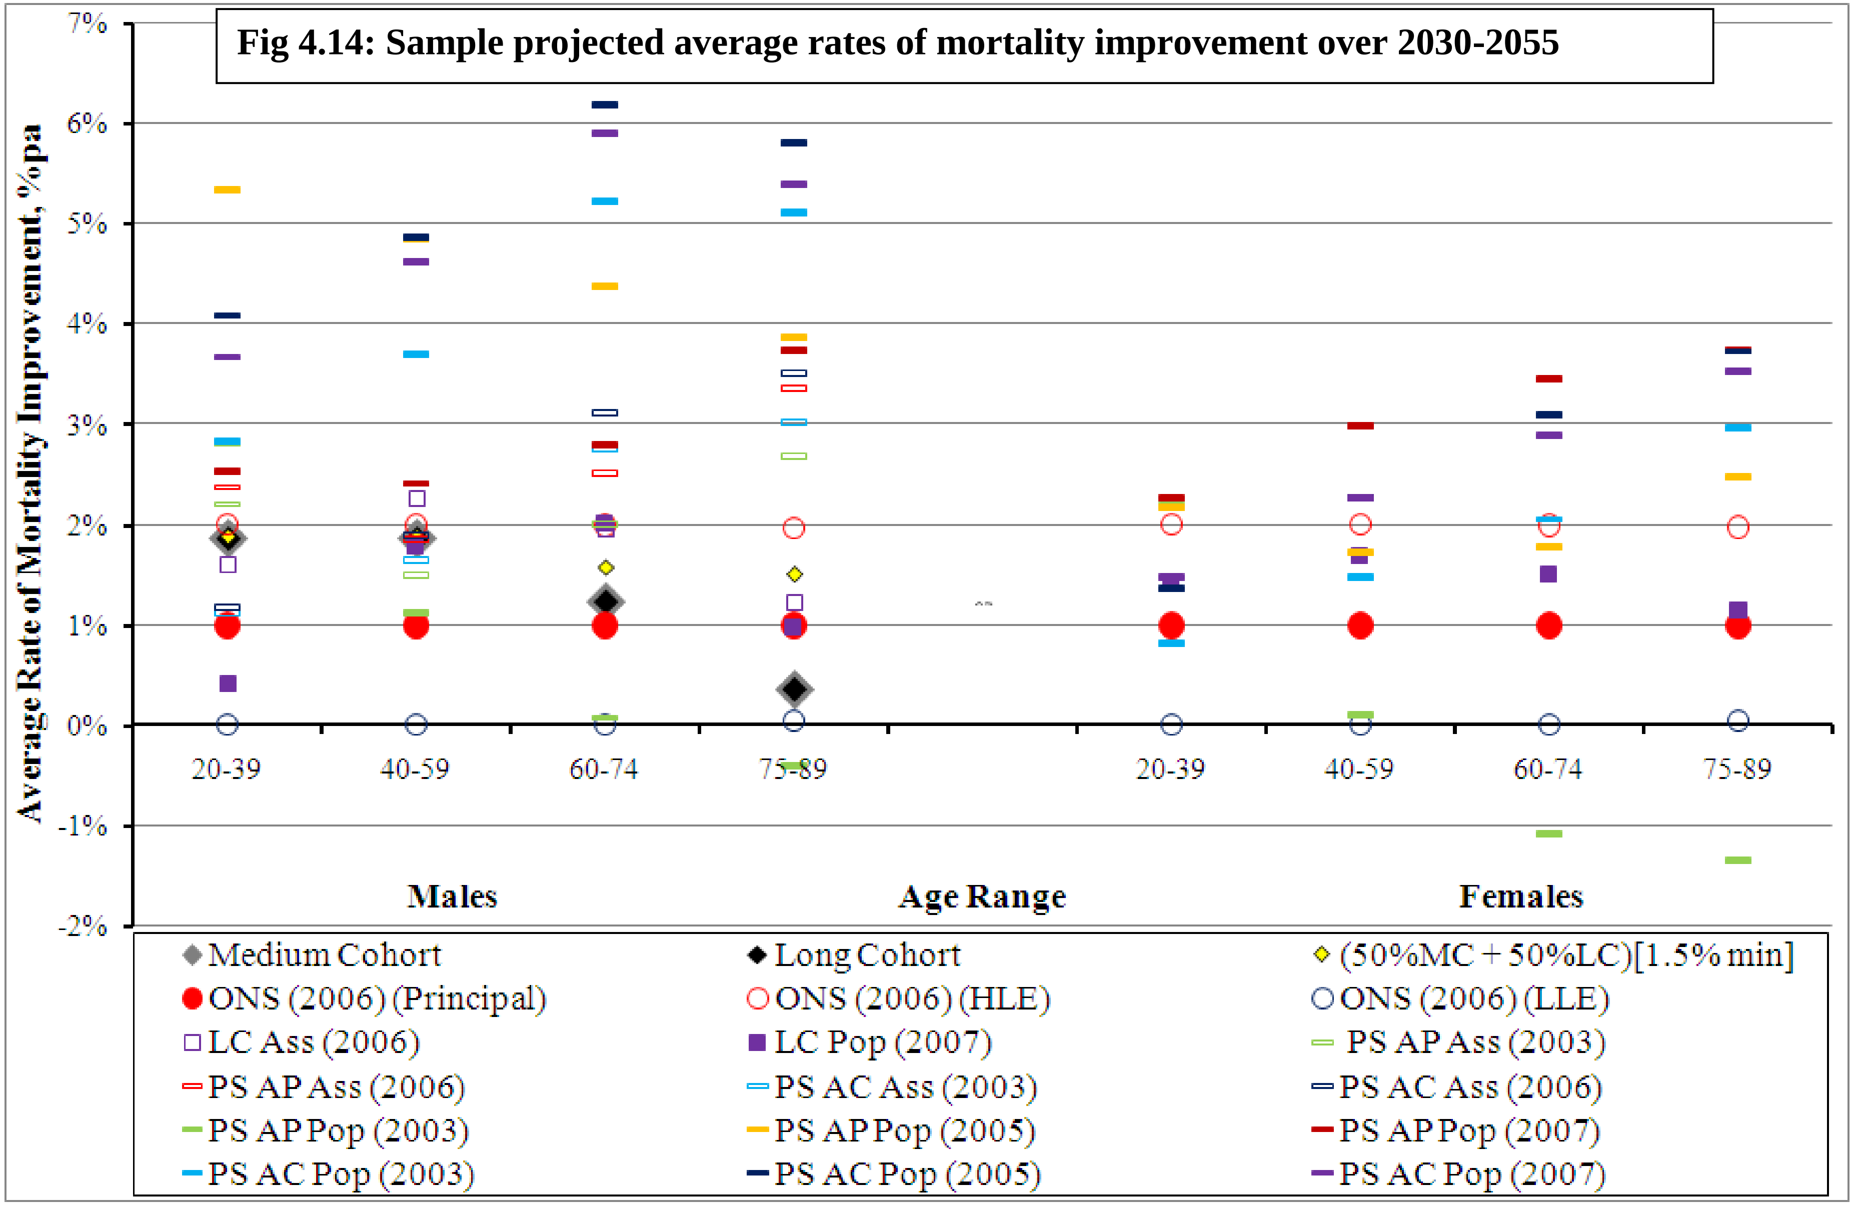

Fig 4.14: Sample projected average rates of mortality improvement over 2030-2055
Average Rate of Mortality Improvement, %pa
7%
6%
5%
4%
3%
2%
1%
0%
-1%
-2%
20-39
40-59
60-74
75-89
20-39
40-59
60-74
75-89
Males
Age Range
Females
Medium Cohort
Long Cohort
(50%MC + 50%LC)[1.5% min]
ONS (2006) (Principal)
ONS (2006) (HLE)
ONS (2006) (LLE)
LC Ass (2006)
LC Pop (2007)
PS AP Ass (2003)
PS AP Ass (2006)
PS AC Ass (2003)
PS AC Ass (2006)
PS AP Pop (2003)
PS AP Pop (2005)
PS AP Pop (2007)
PS AC Pop (2003)
PS AC Pop (2005)
PS AC Pop (2007)

4.3.2. International Perspective on National Population Mortality Projections

The Working Party undertook a small survey of mortality projections produced in various countries. In many cases the methodologies and particularly the assumptions used for the projections are not as readily available as in the UK. In this section we consider the ultimate or long-term rate of mortality improvement within these projections, although it should be noted that this may not be an explicit assumption in some of these projections.

In the UK the National Population Projections produced by GAD – and more recently ONS – have assumed for some time a methodology whereby 'current' rates of mortality change move towards a long-term rate of improvement.

The '2006-based' projections (released in 2007) assumed that annual rates of mortality improvement converged to a common rate of 1.0% p.a. at 2031 for most ages, and continued to improve at this constant rate thereafter. Moreover, it was assumed that those born in the years 1923 to 1940 would continue to experience higher rates of mortality improvement, with assumed rates of improvement in and after 2031 rising from 1.0% p.a. for those born before 1923 to a peak of 2.5% p.a. for those born in 1931 and then declining back to 1.0% p.a. for those born in 1941 or later. In the '2004-based' projections (released in 2005) it was assumed that these cohort differentials would disappear over time with a common rate of improvement of 1.0 % p.a. for all ages by 2029.

The long-term assumption of 1.0% p.a. was believed to be consistent with the average pace of improvement over the whole of the 20th century but appears low compared with rates of change currently being experienced by the population of England & Wales and other UK datasets of pensioners and assured lives.

However, mortality projections made in some countries, notably the US and Canada, have used much lower long-term rates of improvement.

The Society of Actuaries (1998) conducted a survey of experts from different professions in Canada, Mexico and the United States on estimated average annual reductions in US mortality rates by 2020. The mean estimated annual improvement rate for males aged 65 and over was 0.64% among actuaries surveyed, 0.82% among demographers, 0.60% among economists and 0.80% among other experts. Similar average rates of change were estimated for females for the same age range (0.64%, 0.93%, 0.60% and 0.87% for the different groups respectively).

More recently, the 2007 Technical Panel on Assumptions and Methods appointed by the US Social Security Advisory Board, recommended that:
"…assumed ultimate rates of mortality decline by age and sex be increased to an average of 1.00 percent per year to be consistent with those observed during 1953-2003 for the total population."

However, the 2009 Annual Report of the Board of Trustees of the Federal Old-age and Survivors Insurance & Federal Disability Insurance Trust Funds, showed they had adopted less prudent assumptions. It assumed mortality rates for ages 65 and over would decline at average annual rates of about 0.32%, 0.71%, and 1.18% between 2033 and 2083 for three sets of assumptions, respectively designated as 'low cost', 'intermediate' and 'high cost'.

The Trustees commented:
"Experts express a wide range of views on the likely rate of future decline in death rates. For example, the 2007 Technical Panel on Assumptions and Methods appointed by the Social Security Advisory Board believed that ultimate rates of decline in mortality will be higher than the rates of decline assumed for the intermediate projections in this report. Others believe that biological and social factors may slow future rates of decline in mortality."

Mortality projections for the Canadian Social Security Programs have also assumed relatively modest long-term rates of improvement. Recent projections (2008) have projected rates of change in years from 2029 of 0.7% p.a. for ages up to 84, 0.6% p.a. for ages 85-90 and 0.4% p.a. for ages 90 and over.

Projections made in some European countries for annuitants have used more rapid long-term assumptions for mortality change. A paper describing the publication of the German annuity valuation table DAV 2004 R (2005) compares the projected rates of change with those contained in Austrian & Swiss projections. The projected long-term rates of change by age are shown in Table 4.3 below:

Age	Males			Females		
	Germany	Switzerland	Austria	Germany	Switzerland	Austria
20	2.25%	2.19%	1.59%	2.25%	2.19%	2.64%
25	1.81%	2.19%	1.54%	2.23%	2.19%	2.54%
30	1.62%	2.19%	1.49%	2.14%	2.19%	2.43%
35	1.58%	2.19%	1.42%	1.80%	2.19%	2.32%
40	1.57%	2.19%	1.35%	1.48%	2.19%	2.21%
45	1.56%	2.85%	1.28%	1.45%	2.19%	2.08%
50	1.50%	3.15%	1.19%	1.55%	2.19%	1.96%
55	1.47%	3.11%	1.10%	1.63%	2.46%	1.83%
60	1.42%	3.06%	1.01%	1.61%	2.75%	1.69%
65	1.51%	3.00%	0.91%	1.66%	3.00%	1.55%
70	1.68%	2.91%	0.81%	1.83%	3.17%	1.40%
75	1.65%	2.74%	0.71%	1.93%	3.21%	1.25%
80	1.38%	2.47%	0.61%	1.82%	3.04%	1.09%
85	1.03%	2.14%	0.50%	1.49%	2.57%	0.92%
90	0.75%	1.81%	0.40%	1.08%	1.85%	0.75%
95	0.75%	1.52%	0.30%	0.81%	1.15%	0.57%
100	0.75%	1.27%	0.19%	0.75%	0.68%	0.39%

Table 4.3. Average annual (long-term) rates of mortality improvement in projections produced in Germany, Switzerland & Austria

4.4. Proposals for the Core Parameter Layer of the CMI Mortality Projections Model

4.4.1. Long-Term Rate of Improvement

The choice of a suitable long-term rate is clearly subjective but it is hoped that the research presented in sections 4.1 to 4.3 may help guide the selection of suitable parameter values.

The Appendix to Working Paper 38 contained a description of future scenarios (in terms of different causes of death) consistent with projections using alternative long-term rates of change. The Working Party considers that mortality models that decompose trends into constituent causes (either cause-of-death models or disease-based approaches) may provide users with insights into appropriate future rates of change. Moreover, the views of experts in different fields are likely to be especially valuable and the Working Party is very supportive of attempts to build links with other professions as evidenced, for example, by the 'Joining Forces on Mortality and Longevity' multidisciplinary conference in October 2009.

The Working Party considers the Long-Term Rate assumption to be the single most important parameter for Users to set for each projection. Therefore, the Working Party does not propose to set any default value for this parameter group.

4.4.2. Long-Term Rate of Improvement: Pattern at High Ages

The structure of the Model allows for the Long-Term Rate of Mortality Improvement to be varied by age and to be split between Age/Period and Cohort Components.

Although no default parameter value is proposed for the Long-Term Rate, the Working Party does propose to simplify the inputs for the Core parameter layer such that only a single input value is required.

The proposed mapping to the full parameter set is to make the Long-Term Rate equal to the single input value for all ages up to and including age 90. For higher ages, the Long-Term Rate for each calendar year is assumed to reduce linearly from age 90, reaching zero at age 120, and to be zero for ages above 120.

This pattern at high ages reflects the observation that rates of mortality improvement have generally run off towards zero at high ages, but that the age-point of reaching zero has been increasing and currently appears to be in the range 100 to 110.

Users of the Advanced parameter layer may set the Long-Term Rates of Mortality Improvement at the granular level of individual age (and by component) to reflect whatever level and pattern they view as appropriate.

4.4.3. Long-Term Rate of Improvement: Age/Period and Cohort Components

For the Core parameter layer, the Working Party proposes to assign the whole of the Long-Term Rate of Improvement to the Age/Period Component, and to set the Cohort Component of the Long-Term Rate to zero.

The Initial Rates of Mortality Improvement are structured so that the Cohort Components average zero across all years-of-birth. The Cohort Component of the Long-Term Rate should therefore be zero if it is assumed that the influence of current year-of-birth features on patterns of improvement will dissipate completely over the Period of Convergence.

Under these proposals, the Long-Term Rate parameter for the Core layer simply sets a general level of mortality improvement with no particular features by cohort or age, other than a run-off to zero at high-age.

5. Convergence from Initial to Long-Term Rates of Mortality Improvement

5.1. Features Evident in Historic Patterns of Mortality Improvement

The structure of the prototype Model assumes that 'current' rates of mortality improvement will blend over time into an assumed 'long-term' rate of change. Features in recently observed experience are assumed to be a good guide as to the likely pace of change in mortality rates in the near term. The challenge then is to form a view as to how these features may change over time and (probably) dissipate.

For the longer term, although we expect there to be peaks and troughs in improvement rates, it is assumed we are unable to forecast them in any detail and so the Model focuses attention on a more general long-term rate of mortality improvement.

Features of mortality change over time may be described in terms of their 'footprint' on a two-dimensional grid of mortality improvement rates over age and time. This covers the strength of the feature (the differential of improvement rates from the 'background' level), and its spread over age and time (including whether the feature corresponds to a particular cohort grouped by years-of-birth).

The Working Party studied features evident in historic patterns of mortality improvement. This was a qualitative analysis based on heat maps of estimated annual rates of mortality improvement. Given the variety and complexity of features observed, and the uncertainty inherent in measuring rates of improvement, more formal analysis was beyond the scope of available resources.

The heat maps used for this analysis are too numerous for inclusion in this paper, but the sources of data are summarised below:

- Population data for England & Wales
  - The ONS dataset covering 1961-2007 – as covered in section **Error! Reference source not found.** of this paper
  - A dataset obtained from the Human Mortality Database and containing death counts and life-years exposure for each calendar year from 1841 to 2005, and by single years of age (workable range is ages 3 to 103).
- Population data for a sample of other countries, drawing from existing papers illustrating mortality patterns and typically covering the last four decades of the 20th Century:
  - Two Dimensional Mortality Data: Patterns and Projections; S J Richards et al; 2007
  - Patterns of Mortality Improvement over Age and Time in Developed Countries; Andreev & Vaupel; 2005.

Whilst the picture is naturally complex, the Working Party found it helpful to roughly segment observed features into two broad groups: age/period and cohort, although there is no exact distinction between these. The key characteristics of these groups are set out below:

Age/Period Features

- 'Age/Period features' appear as a 'footprint' covering a group of ages, over a period of years;
- The observed features appear in a wide range of shapes and sizes;

- In addition, pure 'period features' may be observed covering a group of calendar years, over a wide range of ages (perhaps 40+ years). These are typically short-term effects, perhaps 5-year cycles at most; more recently showing as 1-year variations above or below the 'norm' and so are 'smoothed out' by the P-Spline methodology (or indeed other methodologies which smooth over time).

Cohort Features

- 'Cohort features' appear as a 'footprint' covering a group of years-of-birth, over a wide range of ages and calendar years;
- They generally appear to run for 25+ years, but are difficult to observe clearly below age 40;
- Many appear to relate to quite narrow groups of years-of-birth;
- But no other cohort features – at least for England & Wales data – are as clear, or as strong, as the 'cohort effect' centred on year-of-birth 1931.

In summary, the history of past improvements in mortality has been characterised by bursts of particularly rapid (or slow) change. These features have varied enormously in duration but have always affected some age ranges much more than others, simply because the relative importance of different causes of death varies so much by age. Particular examples are: the reduction in mortality rates for younger adults due to fewer deaths from infectious diseases in the first half of the 20th century; the reduction in mortality rates for older adults – especially males – due to fewer heart attacks since the 1980s; the reduction in breast cancer mortality for females following the introduction of screening in the late 1980s; and the increase in mortality due to AIDS for younger males.

Layered on top of these 'Age/Period' features are further 'Cohort' features, in which the pace of change has been consistently more or less rapid for individuals in particular birth cohorts relative to others, after allowing for the existence of 'Age/Period' features. The obvious example is the 1931 'cohort effect'.

5.2. Insights from Analysis of Mortality Trends by Cause of Death

The Working Party consider that the study of mortality trends by cause of death is likely to offer useful insights into features observed in the aggregate experience, and to help inform views on the likely future course of current patterns of mortality improvement.

Analysis of mortality experience by cause of death improves our understanding of the drivers of the observed features, such as peaks and troughs, by age and over time, in rates of mortality improvement.

Understanding the drivers of current improvement rates can inform views on the extent to which observed features are linked to the current age and period or to the cohort of lives grouped by their year(s) of birth. For example, trend drivers such as improved road safety, (most) medical advances, or (new) infectious disease outbreaks, are likely to affect mortality for particular age-groups, over a particular timeframe, and act independently of the 'history' of any group of lives: drivers of this type are likely to create age/period effects. However, other trend drivers, such as smoking and obesity, act through cumulative health effects over long periods of time, and so their impact is dependent upon the shared 'history' of a group of lives: drivers of this type are likely to create cohort features.

Understanding the drivers of features in current rates of improvement can also inform views on the likely future timescale over which the feature may be expected to run.

Furthermore, understanding the drivers of past features of experience helps in forming views on both the aspects noted above by creating an historical framework or 'library' of features against which the current observations may be compared.

The Actuarial Profession sessional meeting paper 'Longevity in the 21st Century' (Willets et al, 2004) presented an analysis of the contribution of each cause of death (appropriately grouped) to the overall average annualised rates of mortality improvement, over the period 1989 to 2001, for the population of England & Wales. A summary of the results, for male lives, is shown in Table 5.1. (The corresponding figures for female lives are contained in Table 5.3).

Contributions to Overall Average Annual Rates of Mortality Improvement, 1989-2001							
Age Group	**Infectious Disease**	**Cancers**	**Circulatory**	**Respiratory**	**Other Health Related**	**Violence & Accidents**	**All Causes**
20-29	0.1%	0.2%	0.0%	0.0%	-0.8%	0.4%	**-0.1%**
30-39	0.3%	0.3%	0.4%	0.1%	-1.0%	-0.3%	**-0.1%**
40-49	0.1%	0.6%	1.0%	-0.1%	-0.8%	0.1%	**0.8%**
50-59	0.0%	0.8%	2.1%	0.0%	-0.2%	0.1%	**2.7%**
60-69	0.0%	0.9%	2.2%	0.1%	0.1%	0.0%	**3.4%**
70-79	0.0%	0.5%	1.6%	0.0%	0.2%	0.0%	**2.2%**
80+	0.0%	0.2%	0.9%	-0.3%	0.1%	0.0%	**0.9%**

Table 5.1: Breakdown of contributions, by cause of death,
to overall average annual rates of mortality improvement, over the period 1989-2001
England and Wales population, males
Source: Longevity in the 21st Century; Willets et al, 2004; Table 2.8a

Table 5.2 shows an equivalent analysis for the experience over recent years, and was derived by comparing mortality experience for the period 2005-07 against that for 2001-03.

Contributions to Overall Average Annual Rates of Mortality Improvement, 2002-2006							
Age Group	**Infectious Disease**	**Cancers**	**Circulatory**	**Respiratory**	**Other Health Related**	**Violence & Accidents**	**All Causes**
20-29	0.1%	0.1%	0.2%	0.2%	0.7%	2.1%	**3.4%**
30-39	0.1%	0.2%	-0.1%	0.2%	-0.2%	0.5%	**0.6%**
40-49	0.0%	0.7%	1.1%	0.0%	-0.2%	0.2%	**1.8%**
50-59	0.0%	0.6%	1.4%	0.0%	-0.4%	-0.1%	**1.5%**
60-69	0.0%	0.8%	2.6%	0.1%	0.0%	0.0%	**3.4%**
70-79	-0.1%	0.7%	2.9%	0.5%	0.2%	0.0%	**4.1%**
80+	-0.2%	0.2%	2.0%	0.4%	0.1%	0.0%	**2.4%**

Table 5.2: Breakdown of contributions, by cause of death,
to overall average annual rates of mortality improvement, over the period 2002-2006
England and Wales population, males

Table 5.1 shows low rates of change in mortality for the younger age groups during the 1990s, and high rates of mortality improvement for the age group 50-79, with reductions in mortality from circulatory disease as the primary driver.

However, Table 5.2 shows that the picture for the youngest age group changed rapidly, with significant mortality improvement rates in the 2000s driven by reductions in deaths from 'other health related' causes, violence and accidents. This suggests that the current high rates of improvement for younger adult males are likely to form a relatively short-duration age/period feature.

Table 5.2 also shows that reductions in mortality rates due to circulatory diseases continue to be the strongest drivers of mortality improvements for males above age 40. Indeed the feature appears to have strengthened since the turn of the millennium and its peak effect has moved further up the age range. This suggests at least an element of 'cohort' and a relatively long time frame for this feature.

Table 5.3 shows the corresponding analysis to Table 5.1 for females. In contrast to males, there were positive rates of improvement over the period 1989-2001 at all age groups, except 80+, which was zero. However the peak rate for females (2.8% p.a. at ages 60-69) is much lower than that for males (3.4% p.a., also at ages 60-69).

Contributions to Overall Average Annual Rates of Mortality Improvement, 1989-2001							
Age Group	**Infectious Disease**	**Cancers**	**Circulatory**	**Respiratory**	**Other Health Related**	**Violence & Accidents**	**All Causes**
20-29	-0.1%	0.3%	-0.1%	0.0%	-0.3%	0.5%	**0.4%**
30-39	-0.1%	1.1%	0.0%	-0.1%	-0.4%	0.2%	**0.7%**
40-49	0.0%	1.3%	0.1%	-0.1%	-0.5%	0.1%	**0.8%**
50-59	0.0%	1.1%	1.0%	0.0%	0.0%	0.1%	**2.1%**
60-69	0.0%	0.9%	1.7%	0.0%	0.2%	0.1%	**2.8%**
70-79	0.0%	0.1%	1.6%	-0.5%	0.2%	0.0%	**1.4%**
80+	0.0%	0.0%	0.9%	-0.6%	-0.3%	0.0%	**0.0%**

Table 5.3: Breakdown of contributions, by cause of death,
to overall average annual rates of mortality improvement, over the period 1989-2001
England and Wales population, females
Source: Longevity in the 21st Century; Willets et al, 2004; Table 2.8b

Table 5.4 shows the corresponding analysis to Table 5.2 for females. In almost all age groups the all-causes rate of improvement over the period 2002-2006 has been faster than over the earlier period. In most age groups, the all-causes improvements have been slightly lower than the corresponding figures for males.

The more recent improvements have arisen for a range of causes. Improvements at the younger ages have arisen in most of the causes shown, and most significantly for violence and accidents. At the older ages, reductions in deaths from circulatory disease are the prime driver.

Contributions to Overall Average Annual Rates of Mortality Improvement, 2002-2006							
Age Group	**Infectious Disease**	**Cancers**	**Circulatory**	**Respiratory**	**Other Health Related**	**Violence & Accidents**	**All Causes**
20-29	0.2%	0.4%	0.2%	0.1%	0.2%	0.6%	**1.7%**
30-39	-0.1%	0.6%	0.3%	0.2%	0.1%	0.5%	**1.6%**
40-49	0.0%	1.0%	0.7%	0.1%	0.0%	0.0%	**1.7%**
50-59	0.0%	0.6%	0.8%	0.1%	-0.2%	0.0%	**1.3%**
60-69	0.0%	0.5%	2.1%	0.2%	0.2%	0.0%	**2.9%**
70-79	-0.1%	0.4%	2.6%	0.3%	0.1%	0.0%	**3.4%**
80+	-0.2%	0.0%	1.8%	0.1%	0.1%	0.0%	**1.8%**

Table 5.4: Breakdown of contributions, by cause of death, to overall average annual rates of mortality improvement, over the period 2002-2006 England and Wales population, females

Additional analysis of past trends in mortality by cause of death is presented in the Appendix to CMI Working Paper 38. That analysis covers mortality experience for male lives aged 60-89, in the population of England & Wales, over the period 1968-2005. Deaths were considered in five specific cause-of-death groups, including a focus on heart disease. As well as aiding an understanding of the drivers of past mortality change for males aged 60-89, the analysis formed a basis for developing cause-of-death scenarios for comparison against mortality projections from the Model or other sources.

The Working Party considers that further, more detailed, work in this area is likely to prove valuable, particularly if supplemented by expert medical opinion and analysis of other health indicators and trends in risk factors.

5.3. Proposals for the Core Parameter Layer of the CMI Mortality Projections Model

5.3.1. Period of Convergence

The Model assumes that Initial Rates of Mortality Improvement converge towards Long-Term Rates of Improvement. The convergence process is operated separately for Age/Period and Cohort components, and the Period of Convergence may be varied by age and year-of-birth respectively for the two components.

Users of the Advanced parameter layer may set the Period of Convergence at this granular level to reflect whatever pattern they view as appropriate.

The Working Party considers that the research undertaken to date does justify some variation in the Period of Convergence, by age and component, but only supports a relatively simple, and somewhat subjective, pattern. Having noted that the sensitivity of projection results to changes across the likely range for these parameters is relatively low, the Working Party concluded that it is reasonable to propose a set of Period of Convergence parameter values as a default under the Core parameter layer.

The proposed default values are shown in Figure 5.1.

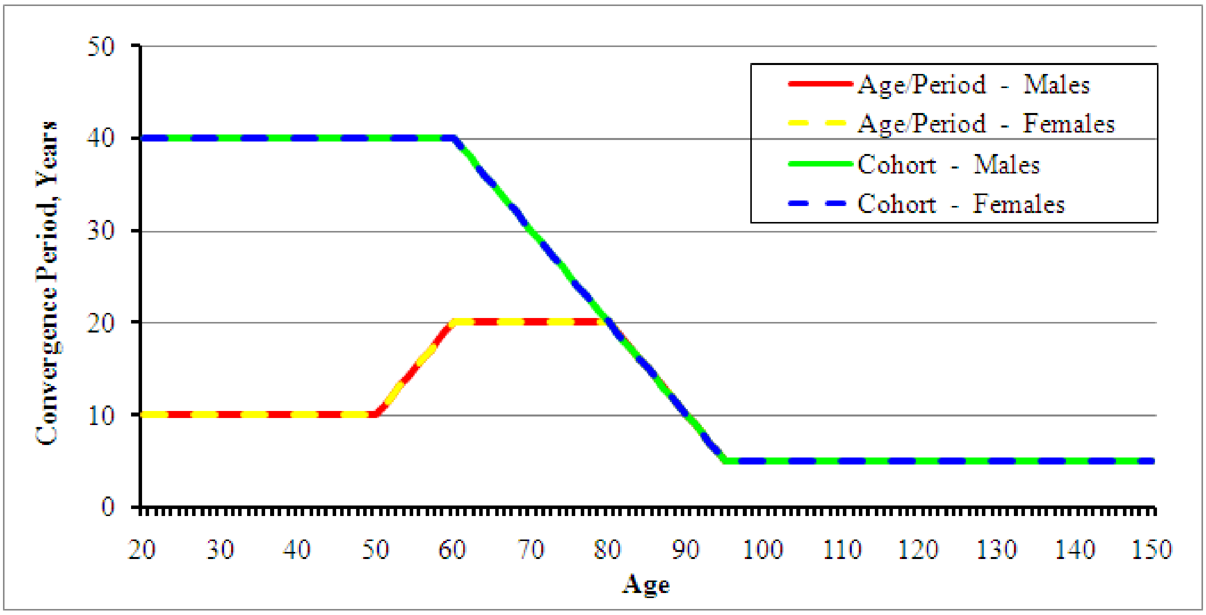

Figure 5.1: Proposed Core parameter layer default values for the Period of Convergence for the Cohort component (x-axis shows age attained in 2005)

The default Period of Convergence proposed for the Age/Period Component is: 10 years for ages up to 50; increasing by one year for each year of age up to 60; then 20 years for all ages to 80; decreasing by one year for each year of age to 95; and then 5 years for ages 95 and above. The same basis is proposed for males and females.

This recognises that Age/Period features in mortality improvements for younger adults have tended to change from decade to decade. However, for ages 60 to 80, where a few major causes of death dominate, there is greater scope for longer-running features and there appears to be no clear end in sight yet for the high rate of mortality improvements driven by advances in relation to circulatory diseases. For the oldest ages it seems appropriate to use a short Period of Convergence, in keeping with the general concept of mortality improvement rates running to zero relatively quickly after age 100.

The default Period of Convergence proposed for the Cohort Component is: 5 years for year-of-birth cohorts 1910 and earlier; increasing by one year for each year-of-birth cohort up to 1945; and then 40 years for all year-of-birth cohorts 1945 or later. The same basis is proposed for males and females.

This recognises that cohort features tend to be long-running, but assumes little continuation beyond age 100. The Periods have been capped at 40 years partly due to concern over the problems of recognising cohort features at younger ages, and therefore the possible inappropriateness of projecting a weakly-based cohort feature far into the future.

5.3.2. Proportion of Convergence Remaining at Mid-Point

The Working Party recognise that features noted in past mortality data come in many shapes and sizes, and that there is little evidence to support any one particular pattern of convergence. Therefore, in designing the Model, the Working Party sought to offer a mechanism for convergence which would be flexible and yet relatively simple, whilst offering a smooth transition into the Long-Term Rates.

Flexibility to influence the initial trajectory of the projected mortality improvements over time, has been achieved in the Model by allowing Users to specify the proportion of the change, between the Initial and Long-Term Rates, that remains at the mid-point of the Period of Convergence.

The algorithm for the convergence process operates by calculating a weighted average of the Initial and Long-Term Rates, with the weights varying over the Period of Convergence.

The proportion of the weight to be placed on the Initial Rate is determined, independently for each individual age and year-of-birth, by fitting cubic polynomials in time, f(t), such that:

- $f(0) = 1$ where $t = 0$ at the start (Foundation Year) of the projection;
- $f(T) = 0$ where T is the length of the Period of Convergence (in whole years);
- $f(\frac{1}{2}T) = P$, the Proportion Remaining at the Mid-Point;
- The first derivative of f(t) at T is zero, that is $f'(T) = 0$.

Let $\tau = t \div T$, so that $\tau = 1$ at the end of the Period of Convergence.

and $f(\tau) = a \times \tau^3 + b \times \tau^2 + c \times \tau + d$

The four criteria above lead to four simultaneous equations:

- $d = 1$
- $a + b + c + d = 0$
- $\frac{1}{8}a + \frac{1}{4}b + \frac{1}{2}c + d = P$
- $3a + 2b + c = 0$

These may be solved to give:

- $a = 8P - 2$
- $b = 5 - 16P$
- $c = 8P - 4$
- $d = 1$

This gives full weight to the Initial Rates of Mortality Improvement at the start of the projection, and to the Long-Term Rates at the end of the Period(s) of Convergence. The convergence path is smooth, passing through the selected level at the mid-point of the convergence period and blending smoothly into the Long-Term Rate. However, the Model has no regard for the trend in rates of mortality improvement prior to the start of the projection, so the transition from historic to projected rates may well not be smooth.

The convergence process is operated separately for Age/Period and Cohort components, and the Proportion Remaining at the Mid-Point parameters may be varied by age and year-of-birth respectively for the two components.

Users of the Advanced parameter layer may set the parameters at this granular level to reflect whatever pattern they view as appropriate and a great variety of transitions may be modelled in this way. Figure 5.2 illustrates how the shape of the convergence can be altered in practice (with a 2.0% p.a. current rate converging towards 1.0% p.a. in 40 years' time).

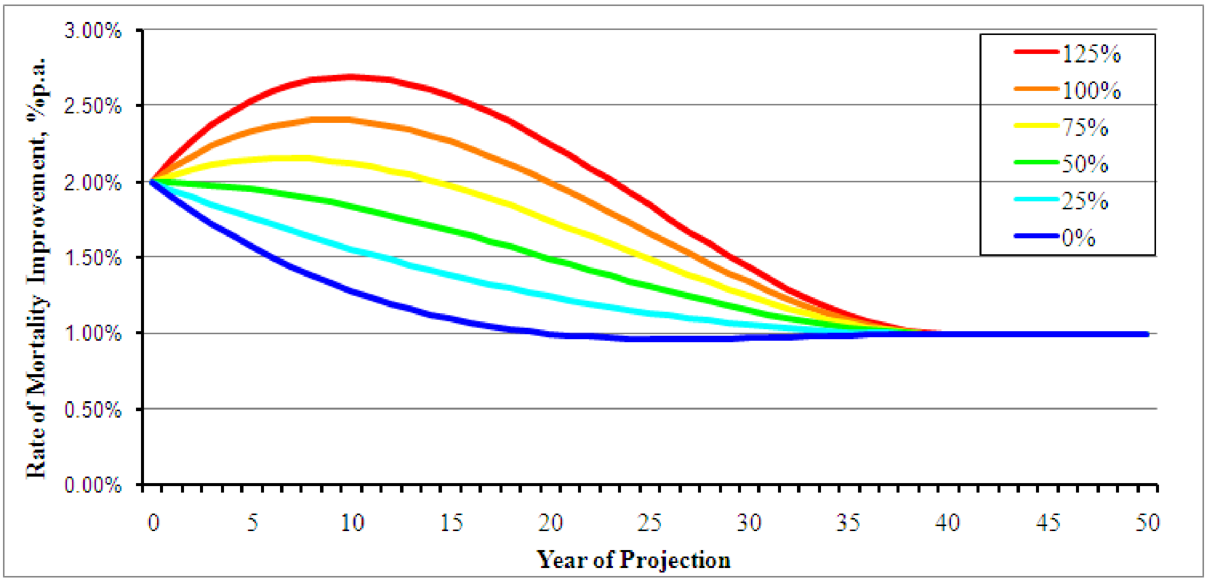

Figure 5.2: Illustration of the operation of the convergence formula over a 40-year period, with various Proportions of Convergence Remaining at Mid-Point

Figure 5.2 demonstrates how selecting a relatively high proportion can generate projected improvements that initially increase before falling towards a long-term average rate. Users therefore have the flexibility to generate scenarios in which the rate of change continues to accelerate in the short-term, before decelerating in the longer term.

Based on the observation that rates of mortality improvement have been increasing in general across a wide range of ages, the Working Party considered setting default parameter values at say 75%. However, recognizing the subjectivity of such an assumption, the Working Party proposes instead to set a simpler, neutral convergence trajectory as a default under the Core parameter layer, allowing Users to take a more active decision on these parameters if they wish.

For the Core parameter level the proposed Proportion of Convergence Remaining by the Mid-Point is defaulted to 50% for all ages and years-of-birth. Therefore, the projected rate of improvement half-way through the Period of Convergence will be the average of the relevant Initial and Long-Term Rates (for both Age/Period and Cohort components).

6. Sample Projections and Parameter Sensitivities

6.1. Sample Projections

Section 4 of CMI Working Paper 38 – 'A Prototype Mortality Projections Model: Part One – An Outline of the Proposed Approach' – illustrated sample Core Projections from the Model by comparing life expectancy and annuity values, for age 65 as at 01/07/2009, for the sample projections against the equivalent values calculated using selected published projections.

An additional perspective on the projections may be gained through the following presentation using the heat map output from the Model to illustrate:

- The pattern of annual rates of mortality improvement by age and calendar year; and
- The projected cumulative change in mortality reduction factors by age and calendar year.

First, the following heat maps show the pattern of mortality improvement rates generated by two sample Core Projections and two published projections:

- Figure 6.1: CPMv0.0 [1.0%] {male}
- Figure 6.2: CPMv0.0 [3.0%] {male}
- Figure 6.3: Medium Cohort_1.0% minimum
- Figure 6.4: 120% Long Cohort_2.5% minimum.

All the projections are for male lives and the projection basis applies from 2006. In all cases the estimated past rates, covering the period 1991 to 2005, are those adopted as the default values under the 'Core' parameter layer of the Model and were derived from population data for England & Wales using an age-cohort P-Spline model.

These charts illustrate the different patterns of projected improvement rates and how the projection fits against estimates of past experience. The examples clearly demonstrate that the Interim Cohort Projections, which were based on experience data to 1999, do not provide a good fit with experience emerging at population level in England & Wales.

In particular, the main cohort feature in the Interim Cohort Projections (centred on 1926) is out of line with that observed in the population (centred on 1931). This difference was recognised in CMI Working Paper 1 (published in 2002), which introduced the Interim Cohort Projections, and noted the 1931-cohort feature in population mortality, but based the Projections around a 1926-cohort feature observed at that time in the major CMI datasets. However, the experience data which has emerged since then is difficult to interpret for insured and pensioner datasets but certainly no longer supports a 1926-cohort feature (see section 3). The charts also confirm visually that the Interim Cohort Projections made no attempt to recognise other patterns in improvement rates, for example at younger ages.

Two further heat maps (Figures 6.5 and 6.6) show the pattern of projected cumulative falls in mortality (measured from the year 2000) for the same two sample Core Projections. By 2060, projected mortality rates are below 50% of the year 2000 mortality rates for a significant range of ages for the Core Projection with a 1.0% p.a. Long-Term Rate, and below 20% for a 3.0% p.a. Long-Term Rate. Although forming a judgment on any form of long-term rate is complex, these charts help to raise the question as to whether it would be useful to allow users to vary the long-term rate over time – for example, in the very long-term the rate of change could be allowed to approach zero, or the mortality rate could 'target' an 'ultimate' level. (See also the Consultation Questions in section 5.3 of Working Paper 38).

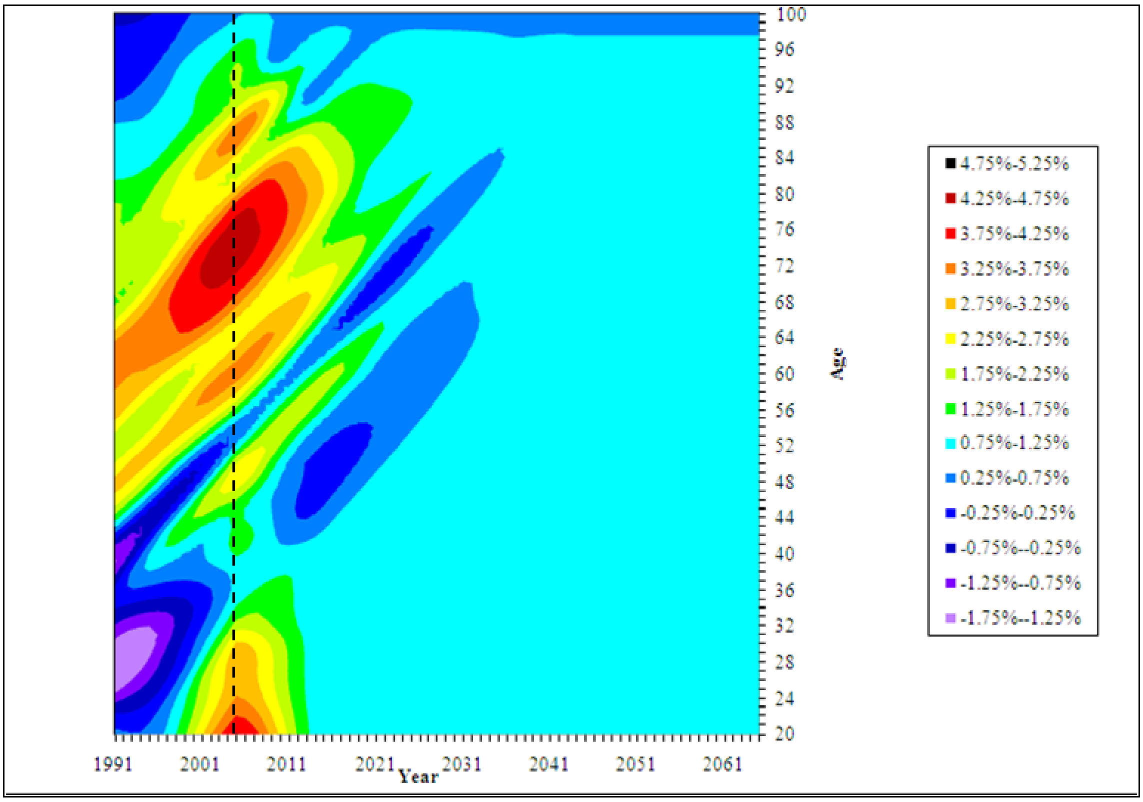

Figure 6.1: Actual and projected annual rates of mortality improvement for males: 1991-2005 – estimated actual rates for population of England & Wales; 2006 onwards – projected rates using CPMv0.0 [1.0%] {male}.

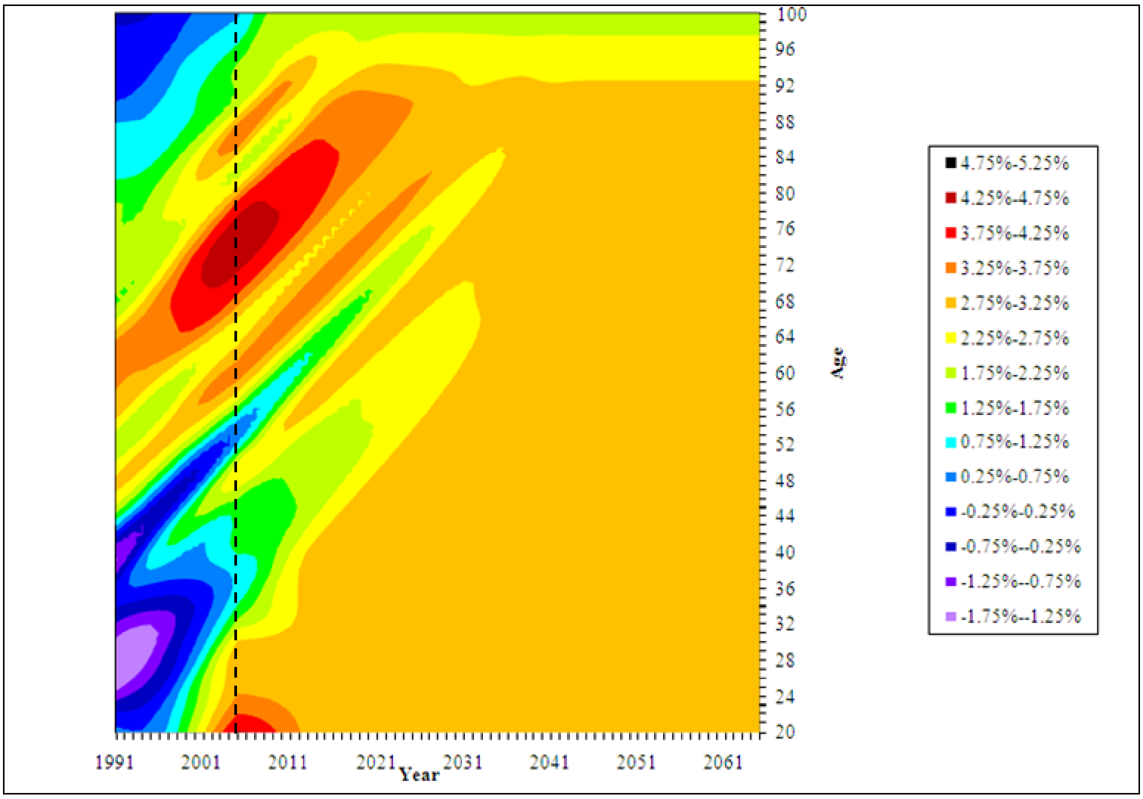

Figure 6.2: Actual and projected annual rates of mortality improvement for males: 1991-2005 – estimated actual rates for population of England & Wales; 2006 onwards – projected rates using CPMv0.0 [3.0%] {male}.

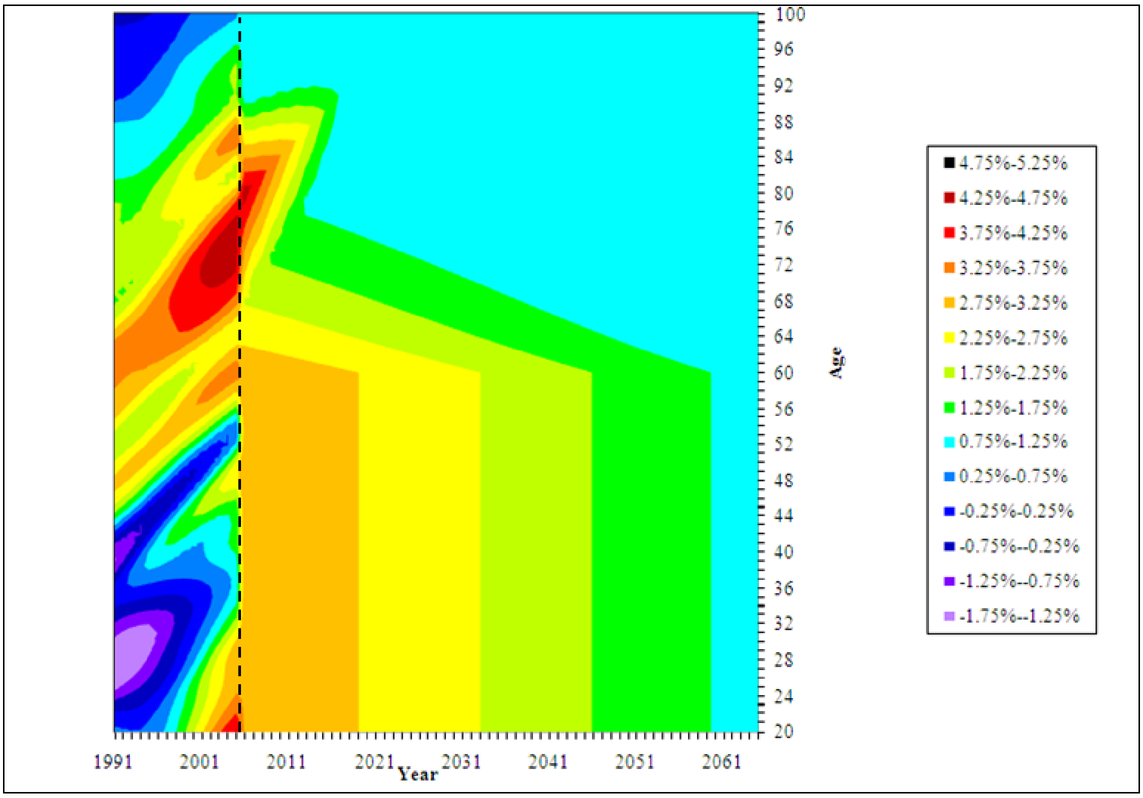

Figure 6.3: Actual and projected annual rates of mortality improvement for males: 1991-2005 – estimated actual rates for population of England & Wales; 2006 onwards – projected rates using 'Medium Cohort_1.0% minimum'.

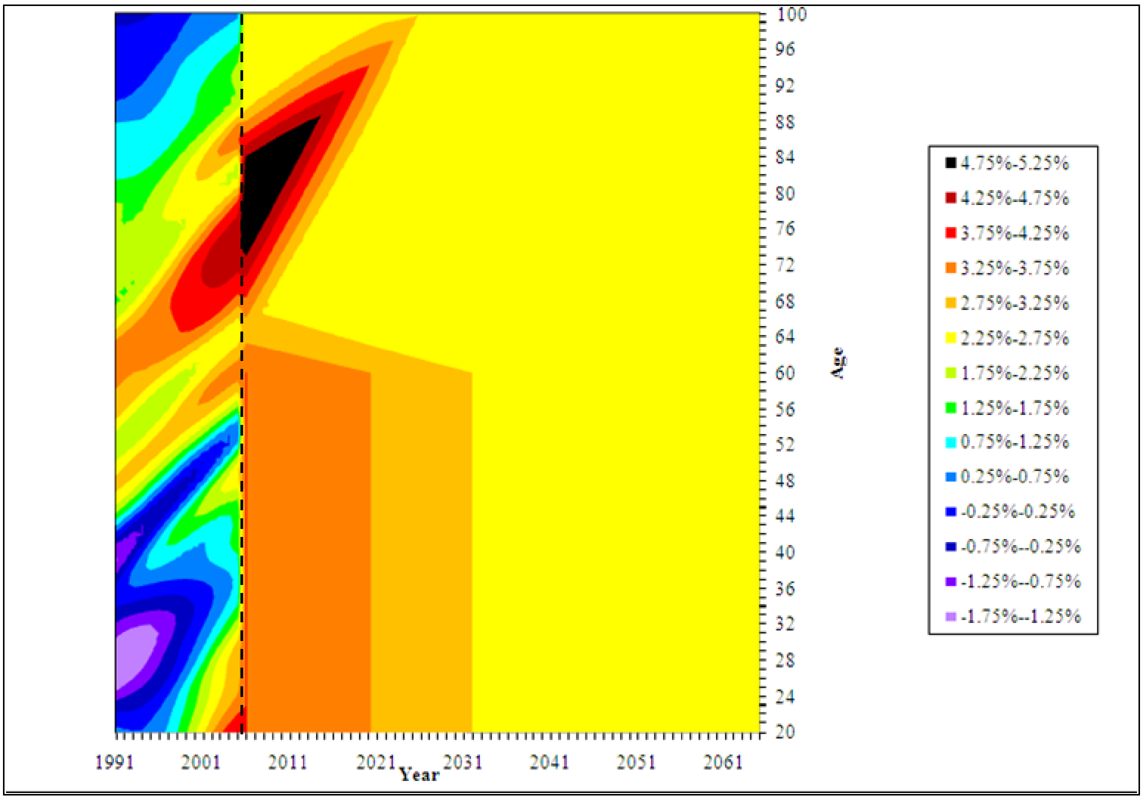

Figure 6.4: Actual and projected annual rates of mortality improvement for males: 1991-2005 – estimated actual rates for population of England & Wales; 2006 onwards – projected rates using '120% Long Cohort_2.5% minimum'.

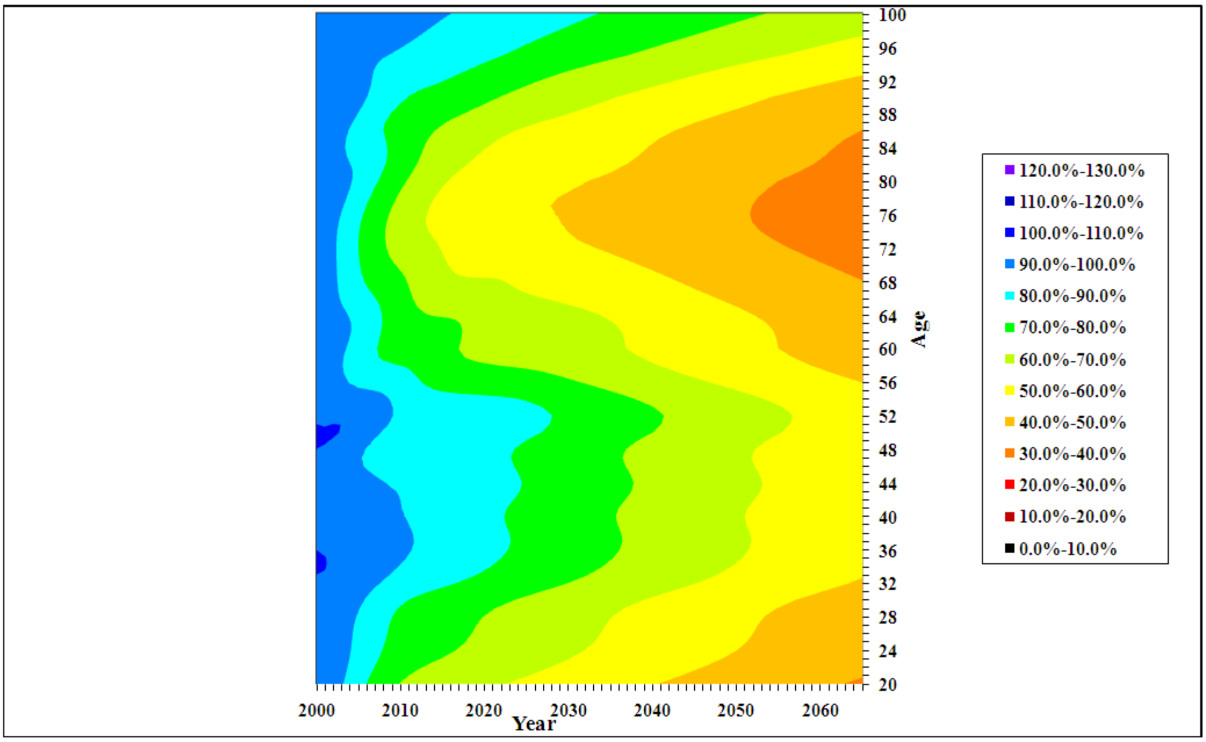

Figure 6.5: Projected cumulative mortality reduction factors for males: 2000-2005 – estimated actual rates for population of England & Wales; 2006 onwards – projected rates using CPMv0.0 [1.0%] {male}.

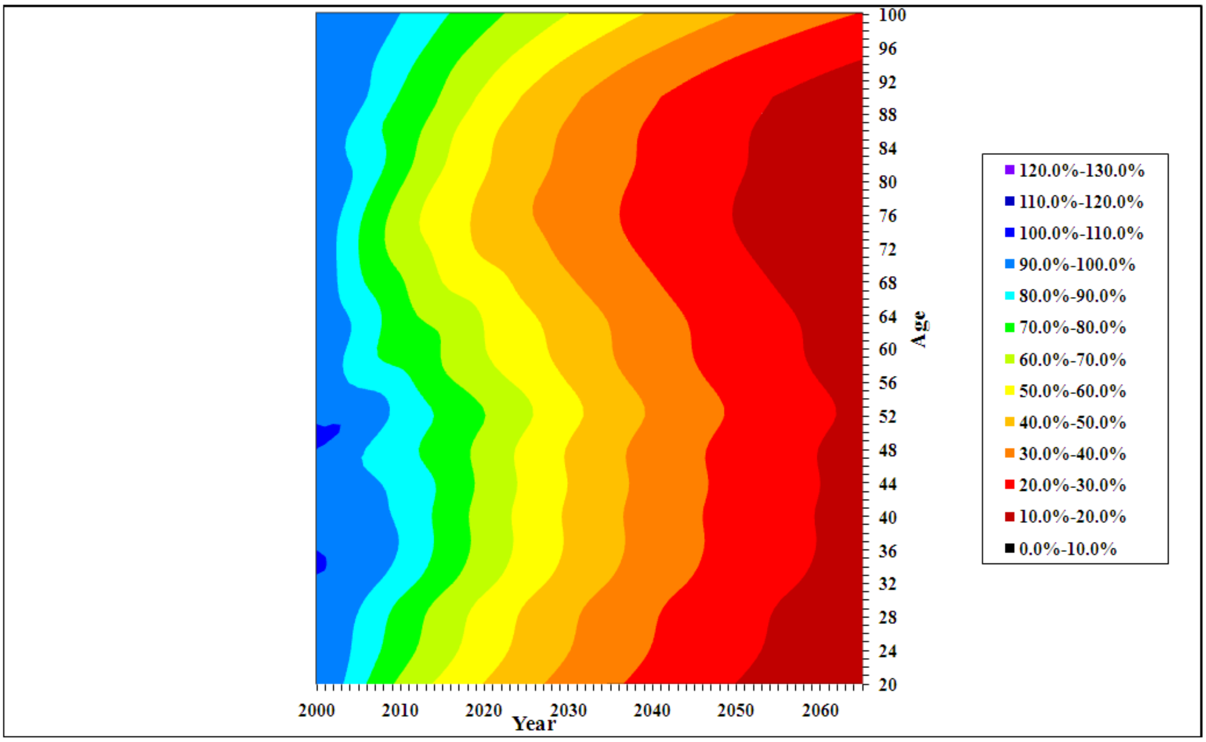

Figure 6.6: Projected cumulative mortality reduction factors for males: 2000-2005 – estimated actual rates for population of England & Wales; 2006 onwards – projected rates using CPMv0.0 [3.0%] {male}.

6.2. Sensitivity of Projections to Changes in the Parameter Values

Sensitivity is illustrated in this paper by reference to values for $\ddot{a}_{65}$ for males. The annuities are assumed to be payable yearly in advance, and values are calculated as at 01/07/2009 using a discount rate of 5.0% p.a. Mortality rates are derived using the PCMA00 base mortality table, projected to 2005 using the past rates of mortality improvement contained in the Model, and forward from 2005 using the rates of mortality improvement given by the various projections listed.

For each of the charts in this section, the different colour bars show the relative values of $\ddot{a}_{65}$ for males calculated on the various projections tested, whilst the 4 blocks of bars show the interaction of these results with changes in the assumed Long-Term Rate of Mortality Improvement. In all cases the result is expressed by showing the resulting annuity value as a percentage of $\ddot{a}_{65}$ for males calculated using the Medium Interim Cohort Projection (from 2005).

Additional results for the sensitivity tests presented are available in a spreadsheet which may be downloaded from the CMI pages on the UK Actuarial Profession's website. Sensitivity is illustrated in the spreadsheet by reference to a selection of annuity values for different ages, and also for the equivalent expectation-of-life values, and for females as well as males; all the annuity and expectation-of-life values are calculated on a cohort, rather than period, life projection basis.

6.2.1. Core Projections

Core Projections are produced when the Core parameter level is selected for all inputs with the user setting values for just two simplified parameters:

- The Long-Term Rate of Mortality Improvement
- A constant Addition to Rates of Mortality Improvement.

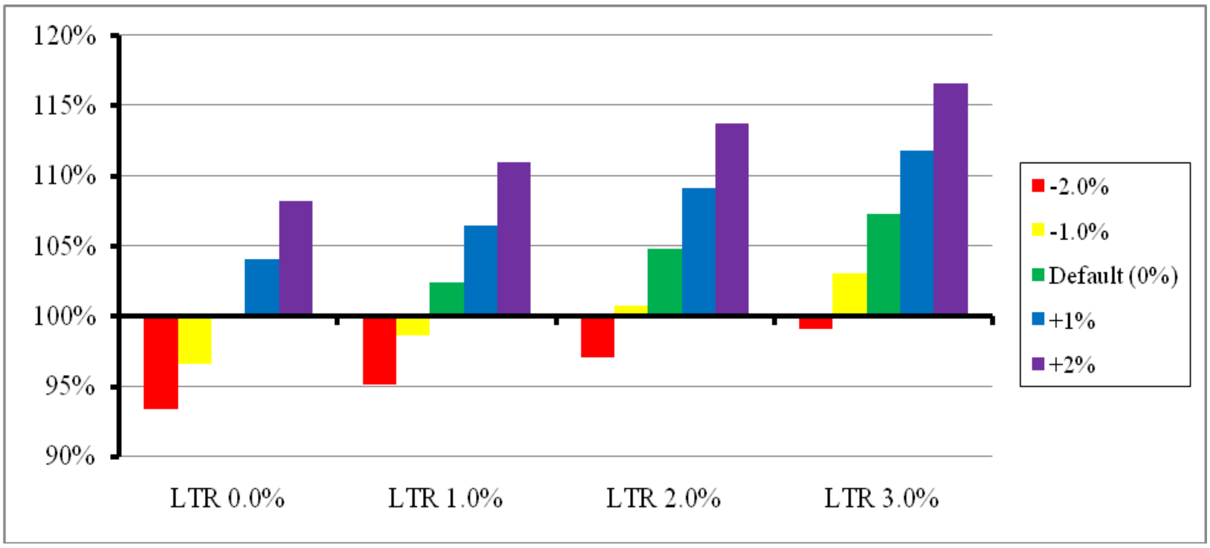

**Figure 6.7: Variation in $\ddot{a}_{65}$ for males,
as % of the value resulting from the Medium Interim Cohort Projection,
for changes in the Constant Addition to Rates of Mortality Improvement
and in the assumed Long-Term Rate of Mortality Improvement**

Figure 6.7 shows the sensitivity of Core Projection results to these two parameters. The $\ddot{a}_{65}$ values for males increase by approximately 2% for each 1.0% added to the Long-Term Rate of Mortality Improvement, and by 4% to 5% for each 1% increase in the Constant Addition to Rates of Mortality Improvement.

The two parameters used for Core Projections represent the most critical inputs to the Model. This is confirmed by comparing the sensitivity of projection results to these two parameters, as shown above, with the sensitivity of results to changes to values set in the extensive Advanced parameter layer, as illustrated in the following sections.

6.2.2. Initial Rates of Mortality Improvement

The default values, applied when the Core Parameter level is selected for this group of inputs, are a set of estimated mortality improvement rates for 2005 derived using an Age-Cohort P-Spline model applied to England & Wales population mortality data. The split between Age/Period and Cohort components is estimated using a relatively simple age-period-cohort model applied to the mortality improvement rates derived from the fitted P-Spline surface. See section 3 for further details.

The sensitivity tests explored are as follows:

- Deriving an alternative estimate of mortality improvement rates for 2005 by using an age-period P-Spline model.
- All Cohort components set to zero – this sensitivity test is consistent with the scenario that cohort features will not have an impact on future mortality improvements and the pattern of future mortality change will be determined solely by age and period elements.
- All Age/Period components set to zero – this sensitivity test is consistent with the scenario that the pattern of future mortality change will be determined solely by year-of-birth cohort elements.
- Maintaining the Cohort component at its current level for ever, rather than the default assumption of running it to a long-term rate of zero for all year-of-birth cohorts.

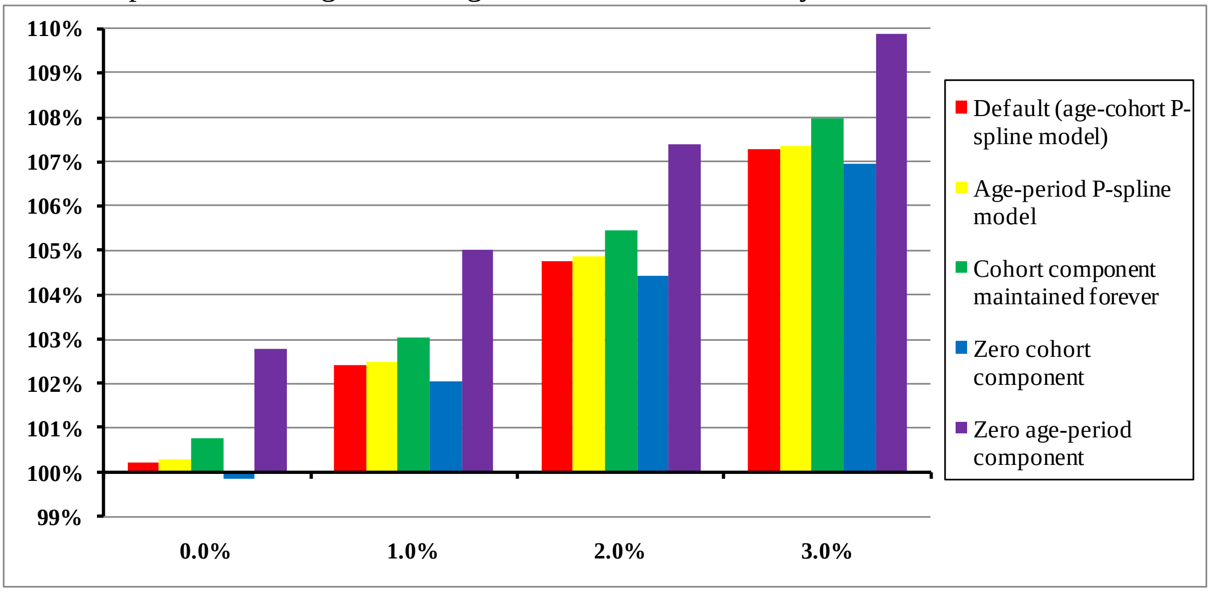

Figure 6.8: Variation in $\ddot{a}_{65}$ for males,
as % of the value resulting from the Medium Interim Cohort Projection,
for changes in the estimated Initial Rates of Mortality Improvement
and in the assumed Long-Term Rate of Mortality Improvement

Figure 6.8 shows that the selected annuity value is relatively insensitive to the choice of age-cohort or age-period P-Spline model for estimating the initial rates of mortality improvement. This conclusion generally holds in aggregate for a pool of business spread across a range of current ages, but more significant differences may arise at specific ages.

However, Figure 6.8 also shows that the results are more sensitive to the split of Initial Mortality Improvement Rates between Age/Period and Cohort components, although the scenarios tested do represent the two extremes of the parameter range. However, even across this pair of scenarios, the change in annuity values is only similar in magnitude to that generated by a 1.0% addition to the Long-Term Rate of Mortality Improvement. Note also that the results for these two particular sensitivity tests are highly dependent on the pattern of initial values by age.

Finally, Figure 6.8 shows that the selected annuity value is relatively insensitive to whether the Cohort component is assumed to run off to zero or to be maintained as a long-term rate (in the context of the other default parameter values).

6.2.3. Long-Term Rates of Mortality Improvement at High Ages

The default mapping, applied when the Core Parameter level is selected for the Long-Term Rate, is for the user-input Long-Term Rate to apply to all ages up to 90, and then for the Long-Term Rate of Mortality Improvement to run, linearly with age, from its assumed value at age 90 to zero at age 120, and to be zero for ages above 120.

The sensitivity tests (alternative mappings) explored are as follows:

- Long-Term Rate drops from its assumed value at age 90 to zero for ages 90 and above.
- Long-Term Rate runs linearly from its assumed value at age 90 to zero at age 150.
- Long-Term Rate holds its assumed value for all ages.

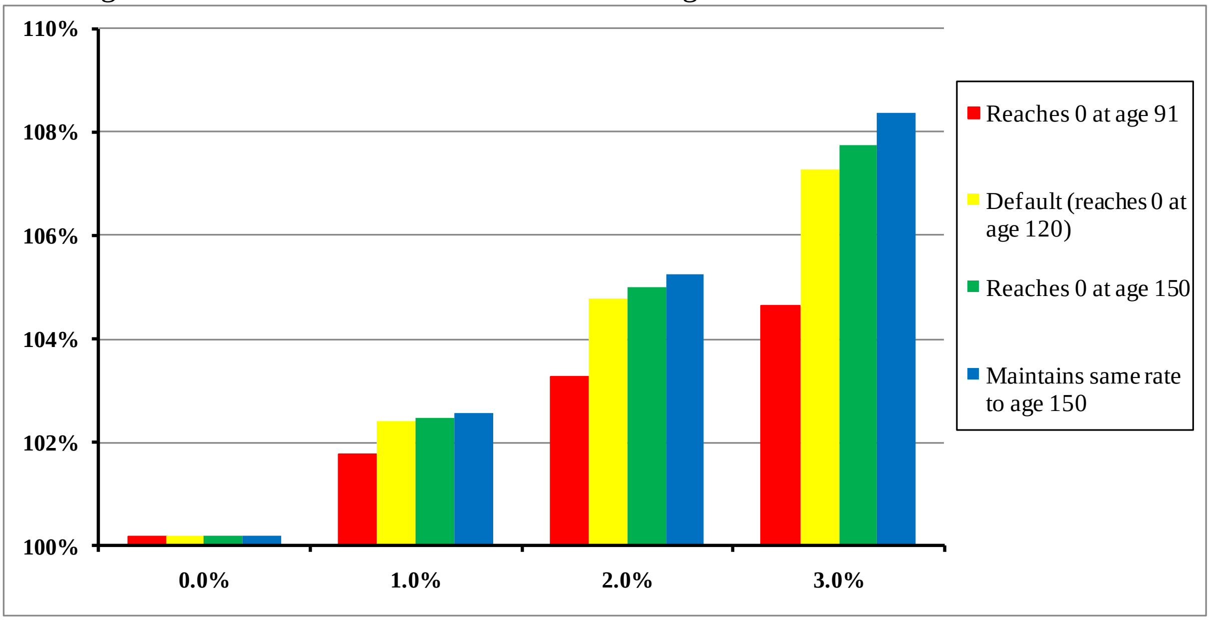

Figure 6.9: Variation in $\ddot{a}_{65}$ for males,
as % of the value resulting from the Medium Interim Cohort Projection,
for changes in the assumed Long-Term Rate of Mortality Improvement
and its pattern at high ages

Figure 6.9 shows that the results are relatively insensitive to the assumed pattern of long-term rates of mortality improvement at high ages. These tests are set in the broad context of assuming that improvement rates are likely to run to zero at some high age, and therefore the results are very dependent on the assumed value for the Long-Term Rate below age 90.

For Core Projections, the default mapping applies the total Long-Term Rate to the Age/Period component and sets the Cohort component to zero – that is it assumes the influence of year-of-birth features dissipates completely over the Period of Convergence. See section 6.2.2 for the results of a sensitivity test in which the Cohort component of the Long-Term Rates is set equal to the corresponding component of the Initial Rates of Mortality Improvement – that is a scenario assuming year-of-birth features continue on indefinitely.

6.2.4. Period of Convergence from Initial to Long-Term Rates of Mortality Improvement

The default values, applied when the Core Parameter level is selected for the Convergence parameter group, varies by age and year-of-birth: see section 5 for details.

The sensitivity tests explored cover increasing Periods of Convergence by 10 years, at all ages, and reducing them by 10 years (subject to a minimum period of 5 years). The tests are applied separately to Age/Period and Cohort components of mortality improvement.

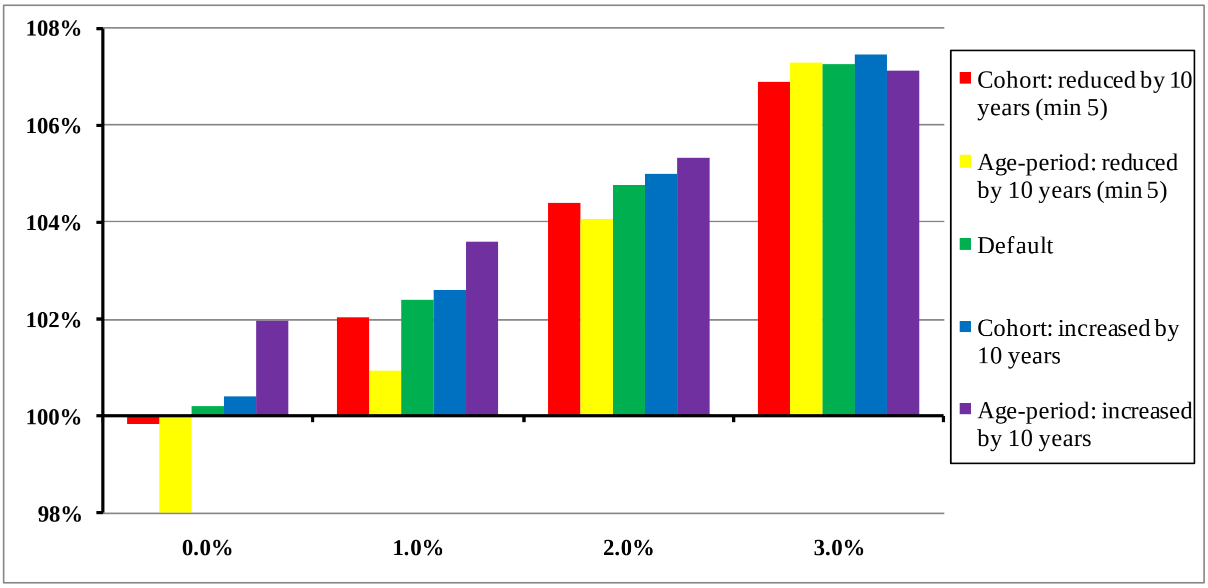

Figure 6.10: Variation in $ä_{65}$ for males, as % of the value resulting from the Medium Interim Cohort Projection, for changes in the Period of Convergence and in the assumed Long-Term Rate of Mortality Improvement

Figure 6.10 shows that the results are relatively insensitive to changes in the assumed Periods of Convergence. Naturally, the sensitivity increases with the size of the gap between Initial and Long-Term Rates, as reflected by the greater variation in results in the left-hand block (Long-Term Rate = 0%) than in the right-hand block (Long-Term Rate = 3%). For the same reason the variation in the left-hand block is much more marked for the Age/Period component (for which the Core parameter defaults give 2.92% for the Initial Rate converging to the assumed Long-Term Rate) than the Cohort Component (for which the Core parameter defaults give -0.43% for the Initial Rate converging to zero, regardless of the assumed Long-Term Rate).

However, even for the Age/Period scenarios, the change in annuity values is generally of lesser magnitude than that generated by a 1.0% addition to the Long-Term Rate of Mortality Improvement until the Long-Term Rate assumption is reduced below 1%.

6.2.5. Pattern of Convergence from Initial to Long-Term Rates of Mortality Improvement

The default pattern, applied when the Core Parameter level is selected for the Convergence parameter group, is that 50% of the change from Initial to Long-Term Rates remains to be achieved after the mid-point of the Periods of Convergence.

The sensitivity tests explored are as follows:

- Proportion Remaining at mid-point is zero – that is the mid-point value is equal to the assumed Long-Term Rate, so that the projected rate converges rapidly to reach its target by the mid-point (although then over-shooting to some degree before regaining its target).
- Proportion Remaining at mid-point is 100% – that is the mid-point value is equal to the Initial Rate, so that the projected rate moves away from its target initially, and then converges to the Long-Term Rate over the second half of the Period of Convergence. For example, this could reflect a scenario where rates of improvement have been increasing and are projected to increase further in the short-term before falling back to a lower long-term rate.
- Proportion Remaining at mid-point is 125% – that is, a more extreme version of the '100% scenario', in which the rate diverges initially and is further away from its target at the mid-point than at the beginning.

In all cases the pattern illustrated is applied to both Age/Period and Cohort components.

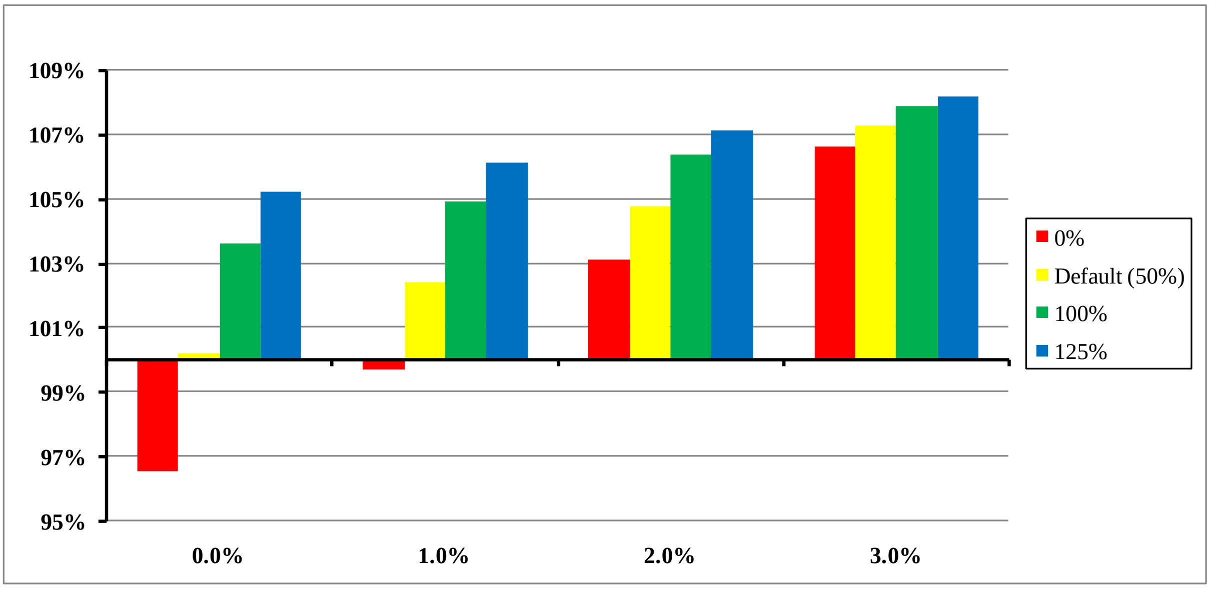

Figure 6.11: Variation in $\ddot{a}_{65}$ for males,
as % of the value resulting from the Medium Interim Cohort Projection,
for changes in the pattern of Convergence
and in the assumed Long-Term Rate of Mortality Improvement

Figure 6.11 shows the results are somewhat more sensitive to the pattern of convergence than to changes in the Period of Convergence, although the scenarios tested may be thought of as being towards the extremes of realistic. In particular, given these results and the general trend of increasing mortality improvement rates across a wide age-range over the past 20 years, it seems appropriate for users of the Model to prioritise further consideration of using this parameter to reflect their view of the likely short-term path of rates of mortality improvement, and how this path may blend into convergence to their Long-Term rate assumption over a longer period.

References

“Patterns of Mortality Improvement over Age and Time in Developed Countries.” Andreev & Vaupel (2005).

Annual Report of the Board of Trustees of the Federal Old-age and Survivors Insurance & Federal Disability Insurance Trust Funds (2009)

CMI Working Paper 1: An interim basis for adjusting the "92" Series mortality projections for cohort effects (2002)

CMI Working Paper 15: Projecting future mortality: towards a proposal for a stochastic methodology. (2005)

CMI Working Paper 20: Stochastic Projection Methodologies: Further progress and P-Spline model features, example results and implication (Revised version, 2007)

CMI Working Paper 37: Version 1.1 of the CMI Library of Mortality Projections (2009)

CMI Working Paper 38: A Prototype Mortality Projections Model: Part One: An outline of the proposed approach (2009)

Smoking behaviour and socio-economic class: a cohort analysis, 1974 to 1998. Evandrou, E. & Falkingham, J. Health Statistics Quarterly, 14, 30-38. (2002)

“Mortality Projections for Social Security Programs in Canada.” M. Montambeault & J.-C. Ménard (2008)

ONS Longitudinal Study age-specific mortality data by social class, 1972 – 2005. Office for National Statistics (2007)

“Coping with longevity – the new German annuity valuation table DAV 2004 R.” U. Pasdika & J. Wolff (2005)

“Two-dimensional mortality data: patterns and projections” Richards S. J., Ellam J. R., Hubbard J., Lu J. L. C., Makin S. J. and Miller K. A. (Presented to the Faculty of Actuaries, 19 March 2007 and to the Institute of Actuaries 24 September 2007).

"Summary of Results of Survey of Seminar Attendees." Marjorie Rosenberg and Warren Luckner. *North American Actuarial Journal* 2(4) (October) 75 (1998).

“Report to the Social Security Advisory Board”. Technical Panel on Assumptions and Methods (2007)

“Longevity in the 21st century” R C Willets, A P Gallop, P A Leandro, J L C Lu, A S Macdonald, K A Miller, S J Richards, N Robjohns, J P Ryan and H R Waters, presented to the Faculty of Actuaries on 15 March 2004 and to the Institute of Actuaries on 26 April 2004