ISSN 2044-3145

Continuous Mortality Investigation

High Age Mortality Working Party

WORKING PAPER 85

Initial report on the features of high age mortality

OCTOBER 2015
(This Working Paper was reissued in October 2015 to incorporate a minor amendment to the axis-labelling of Figure 4.1)

Correspondence address: Cheapside House, 138 Cheapside, London, EC2V 6BW
Tel: 020 7776 3820 Fax: 020 7776 3810
www.cmilimited.co.uk info@cmilimited.co.uk

Continuous Mortality Investigation Limited ("CMI") is registered in England & Wales
(Company number: 8373631) with its Registered Office at: Staple Inn Hall, High Holborn, London, WC1V 7QJ

Contents

1. Executive Summary

The purpose of this paper is to provide a report from the CMI High Age Mortality Working Party on its research and findings to date. The Working Party was established in 2014 with the following broad aims (further details are provided in the Terms of Reference listed in Appendix A):

- Provide a broad indication of the potential financial impact of misestimating high age mortality.
- Investigate and summarise published research on high age mortality.
- Identify potential issues with existing data sources used by CMI (Self-Administered Pension Schemes (SAPS), insurers, Office for National Statistics (ONS)) and methodology specific to high age mortality.
- Outline specific analysis to assess the impact of any identified issues and propose potential ways to enable the CMI to address them.

The key issues identified and considered are:

- The expert judgement required to set a high age mortality assumption (in terms of level and shape) where portfolio levels of data are currently low.
- The issues to consider when using population mortality to 'fill this gap'.
- The impact of age misreporting and age at death misstatement on reported mortality.

The key findings from our work to date are:

- There are various issues with the data quality for all the data sources which might be used to model mortality at the oldest ages. It is not immediately clear how much the differences in results might be due to data issues. As such, it is difficult at this point to make any recommendations on a preferred approach or approaches to modelling mortality rates at the oldest ages. However, the differences in the resulting mortality rates derived using different models are generally not material except at the very oldest ages. Hence, it is unlikely that any one approach would produce results which are materially different from another except at the oldest ages. Given this and the data issues mentioned above, the Working Party does not feel that the choices made to extrapolate mortality rates at the oldest ages in recent graduated CMI tables were unreasonable.
- Analyses of historical mortality from extinct cohorts imply that mortality for the England & Wales population above age 90 as published by the ONS has been underestimated by around 5% for males and about 1-2% for females in the period considered. We understand the underestimation of mortality to be driven by overestimation of population exposures at very high ages. The Working Party intend to consider this feature further in the next phase, including the potential impact on both current (base) levels of mortality and on the use of revised estimates on future mortality projections. Please see section 7.1 for further details.
- There is wide variation in the level and shape of mortality assumed at high ages under different tables published by the CMI, the ONS and North American actuarial associations. If the ranges of methodologies typically adopted are applied to the data underlying the S2PML tables then we observe a variation of -4.1% to +0.4% in cohort

life expectancy (or -3.7% to +0.3% in annuity value with a 3% discount rate) for a male aged 90. Impacts are smaller at age 65, being -0.4% to +0.1% on cohort life expectancy and -0.2% to +0.0% on annuity value. Please see section 5.5 for further details.

- The debate on whether the shape of mortality at high ages is exponential or exhibits signs of mortality deceleration is inconclusive. The S2PML tables effectively assume mortality deceleration occurs. The impact on cohort life expectancy for a male aged 90 of instead adopting a Gompertz mortality shape at high ages is -2.5% to -0.5% (or a -2.0% to -0.4% impact on annuity value). Again, impacts are smaller at age 65, being -0.4% to -0.1% on cohort life expectancy and -0.2% to -0.0% on annuity value. Please see section 4.2 for further details.

- We have modelled scenarios considering the impact of late reporting of deaths at high ages. The potential impact on annuity values and life expectancies for a male aged 65 if late reporting is understated is typically in the region of +1% to +2% (and can be as high as +5% to +10% at age 90 in some of the scenarios considered). The delay in reporting deaths has a more material impact at very high ages where the higher rate of mortality will have a bigger impact on restating exposures once all deaths are known. Please see section 6.1.4 for further details.

- We have modelled scenarios considering the impact of age misstatement. This suggests the potential impact on annuity values and life expectancies for a male aged 65 from age misstatement is in the region of 0.0% to +0.1% (and up to +0.5% to +1.0% for a male aged 90). We have considered the impact of date of birth corrections from data cleansing of large buy-in clients of a bulk annuity provider in constructing these scenarios. The impact varies by the range of age misstatements. Please see section 6.1.5 for further details.

For the next phase, the Working Party intends considering the following areas:

- Population mortality at high ages: we intend to consider further the appropriateness of estimated mortality for the England & Wales population. In their 2014 paper, ‘Phantoms Never Die: Living with unreliable mortality data’, Cairns et al indicate issues associated with the ONS re-stating their point estimate for population exposures above 85 across all high ages using extinct generation methodologies. There are also concerns around the mid-year population estimates for particular cohorts, most notably for lives born just after the end of the First World War. We intend to explore these areas further and consider the appropriateness of the ONS approach and whether there are more appropriate alternative approaches. We will also explore whether there are any other datasets, such as longitudinal studies, which may provide additional useful information.

- Mortality trend: this paper focuses on issues that may affect historical and recent levels of mortality. We intend to explore how the analysis described above might help inform a time analysis of mortality trends, in particular for assisting the CMI Mortality Projections Committee with their modelling at high ages.

- International comparators: we also intend to extend our analyses to consider other territories and consider where similar features, or otherwise, are observed relative to the England & Wales population.

Our intention is to produce a follow-up paper considering these themes in 2016.

2. Introduction

This paper reports on the research and findings to date of the CMI High Age Mortality Working Party. It sets out our initial findings on data and modelling approaches along with consideration of potential issues, implications and areas for further investigation. There has been a lot of recent research in this area and this paper seeks to pull together the key areas considered. Our intention is to provide a further update to our work in 2016.

The Working Party was drawn predominantly from members of CMI investigation committees and comprised: Steve Bale (Chair from April 2015), Mark Cooper, Andrew Gaches, Adrian Gallop, Joynur Rahman and Neil Robjohns (Chair until April 2015).

For clarification, we have focused on mortality for ages 90 and above. For the England & Wales population, this represents 0.8% of the population in 2013, and is projected by the Office for National Statistics (ONS) to grow to 2.6% by 2040. The ONS 2012-based population projections forecast the proportion of deaths at ages 90 and above to grow from 20% of all deaths in 2013 to 36% of all deaths in 2040. Our choice to focus on ages above 90 reflects the desire to understand features of very high age mortality given the increasing impact of higher age mortality from an ageing population. For many insurance companies and pension schemes, mortality experience and data coverage will predominantly occur at younger ages, so understanding the uncertainty around the very high ages is critical.

In addition, we have focused on pensions and annuities products; with limited consideration of the issues affecting whole of life products to the extent that these differ.

Note that throughout the paper we provide future life expectancies as a commonly accepted metric of assessing the impact of modelling and data issues.

The subject matter of this paper crosses a number of CMI committees' work. The Working Party wishes to acknowledge and thank Bill Baker, David Bartlett, Deborah Cooper, Mary Hall and Steven Rimmer for their valuable review of a draft version of this paper. The Working Party is also grateful to Club Vita and Rothesay Life for each supplying the Working Party with results from their analysis of pensioner experience.

The structure of this paper is as follows:

- Section 3 sets out the context to this paper, in particular why high age mortality matters;
- Section 4 summarises a review of theories on the shape of mortality at very high ages, in particular whether mortality decelerates or not;
- Section 5 considers the range of techniques adopted for recent mortality tables and their range of financial impacts;
- Section 6 considers features and issues of high age datasets for both population and insured / pension scheme data;
- Section 7 summarises initial findings from exploring the mortality of closed cohorts, including high-level trends;
- Section 8 considers whether mortality rates converge with increasing age;
- Section 9 summarises the key findings from sections 4 to 8; and
- Section 10 lists our intended next steps.

3. Context setting

3.1. Why does High Age Mortality matter?

The understanding of developments in old-age mortality is important for:

- Governments providing social security, health, housing and welfare programmes;
- Private healthcare / care home providers;
- Insurance companies and pension funds providing benefits in retirement;
- Investment product providers for income drawdown; and
- Investment markets providing longevity swaps / solutions.

The ONS estimate in their 2012-based population projections that the number of centenarians (lives aged 100 and above) could increase from just under 14,000 in 2014 to over 250,000 by 2051. This observed and expected growth in centenarians is illustrated in Figure 3.1.

The fact that people are living longer has rightly been widely celebrated. However, an ageing population leads to higher expenditure on pensions and health, and such expenditure could potentially grow faster than tax receipts. The estimation of populations at high ages is therefore important for Governments providing social security, health, welfare and housing programs. The impact on these programmes will be influenced by the change in the proportion of lives requiring social care, along with their expected remaining lifetimes.

Figure 3.1: Number of centenarians – historical and projected, England & Wales

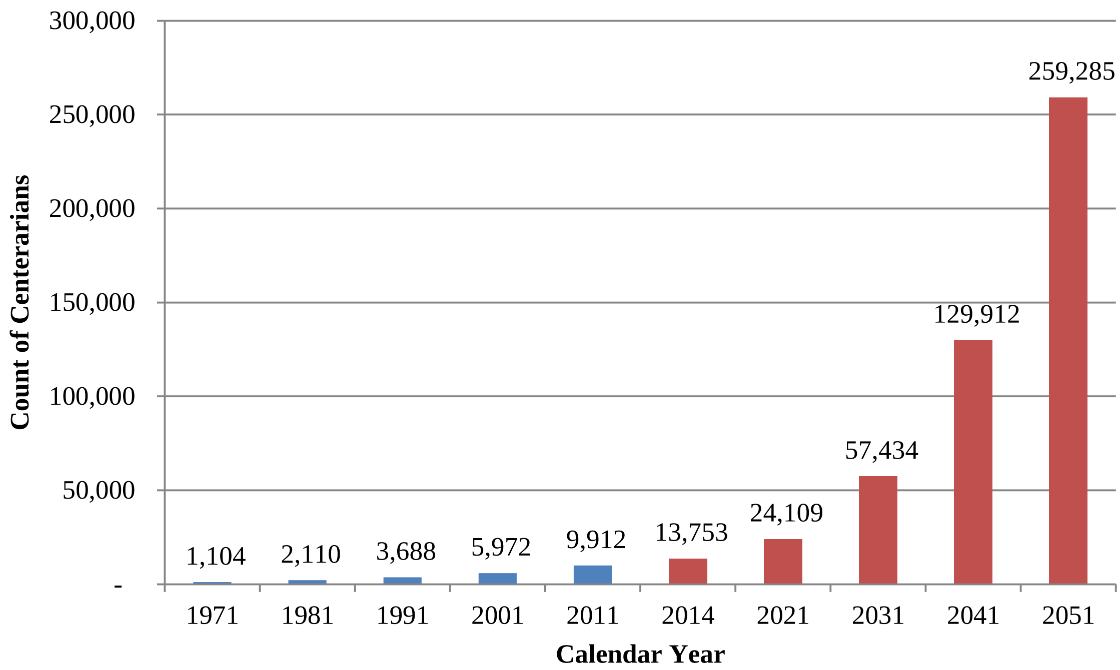

Source: ONS, 2012-based England & Wales principal population projections. Blue bars represent historical numbers; red bars represent future projected numbers.

A better understanding of the uncertainty and fluctuations in estimating old-age mortality will be important for actuaries involved in the pricing of retirement products and for capital allocation and reserving using an internal model under Solvency II. Furthermore, the impact of the 2014 Budget on retirement planning is expected to result in an increased set of retirement income provision options through a combination of asset and insurance solutions. Consumer behaviour might result in a shift in product preferences at different ages.

The availability of data for setting assumptions will vary according to the volumes of data at each age. Shorter duration in-force annuity business may well have a younger average age and assumption-setters may need to consider other available data sources, such as industry and / or population information to determine appropriate assumptions for high ages.

Figure 3.2 shows the proportions of those born today who are expected to survive to old ages based on the survival factors, l_x, taken from the UK National Life Tables 2011-13 with no allowance for changes in future mortality. We can observe that at age 90 whilst about 20% of males are expected to be alive, the expected proportion of females alive is higher at 30%. This illustrates the particular importance of setting mortality rates for very high age females, particularly for survivor retirement benefits (spouse's pensions). The impact on the cost of meeting the liabilities will depend on the age mix of a particular portfolio. Portfolios where there is a high proportion of spouse's pension benefits and/or a high proportion of insured females will be more sensitive to a change in mortality at very high ages, as we illustrate later.

Figure 3.2: Proportion alive at each age – UK National Life Tables 2011 – 2013

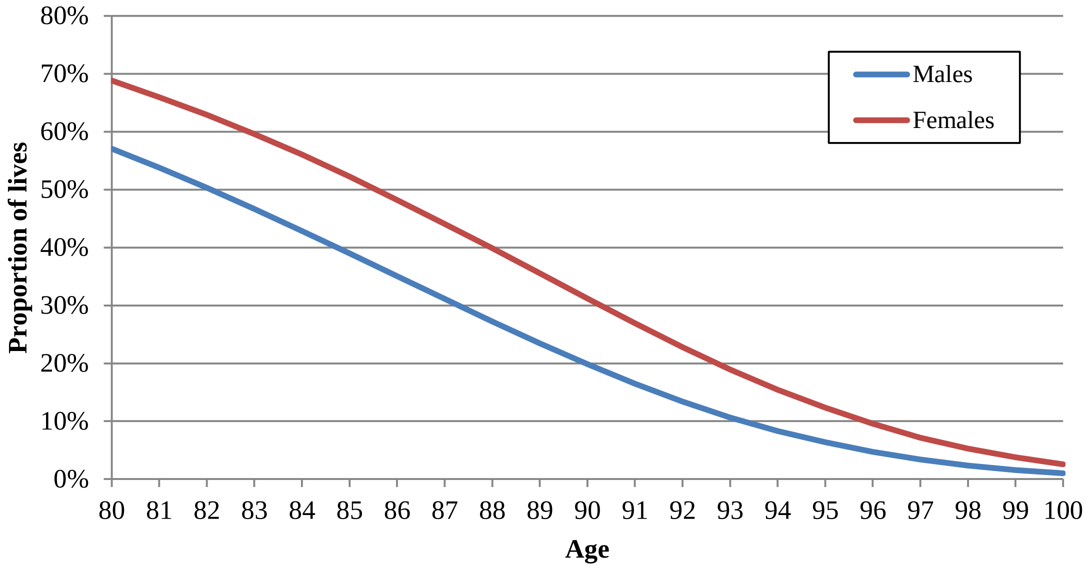

Source: National Life Tables 2011-2013, United Kingdom, ONS

3.2. Data volumes at high ages

The graphs in Figure 3.3 compare the levels of exposure and deaths for the period 2007-10 for the England & Wales population, CMI All Offices Annuities dataset and CMI SAPS populations.

Figure 3.3: Deaths and exposure by age for various mortality sets, 2007-2010

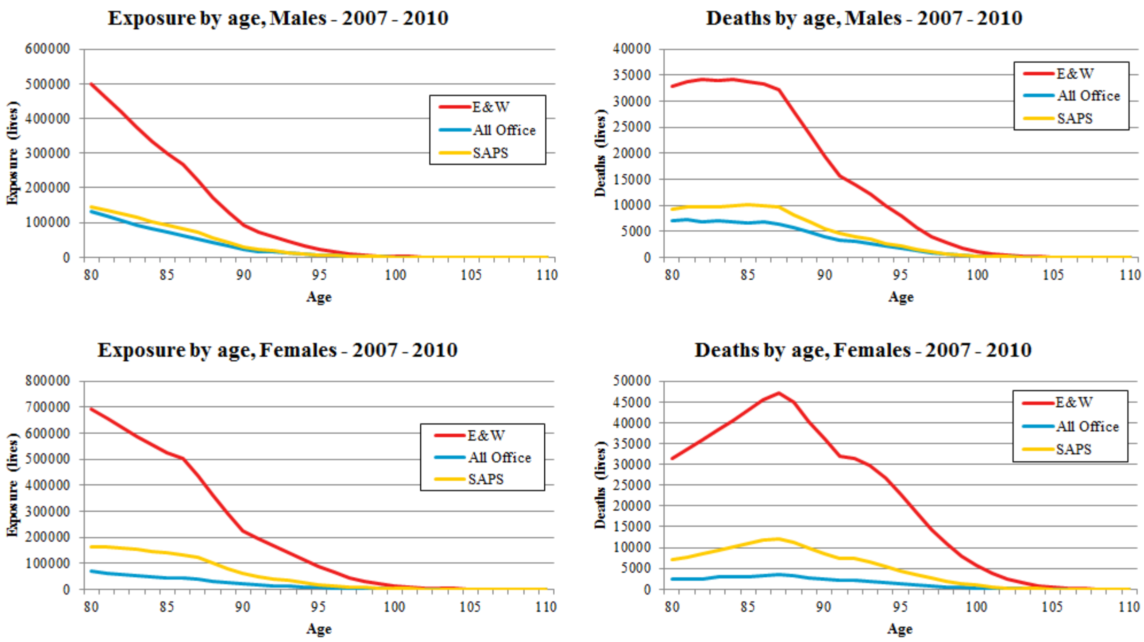

Source: ONS, CMI

We observe that for both genders, the lives exposed reduce rapidly, and to particularly low levels by age 90. A similar feature is observed for deaths.

At higher ages the England & Wales population mortality data is more than two to three times the size of the insured data. Readers may wish to consider the appropriateness of moving from reliance on portfolio or CMI experience to population mortality experience when setting mortality assumptions for age 90+. We consider this further in this paper.

4. Review of competing theories for patterns of high age mortality

4.1. Mortality Deceleration – a discussion

One of the key areas of debate regarding the extension of mortality curves into higher ages is whether the rate of increase of mortality rates "decelerates" away from a Gompertz law (log μ_x is linear in x). Observed mortality rates from most datasets tend to plateau or decrease at older ages but this is generally ascribed to poor data quality as discussed in section 6. However, there remains a debate over whether the Gompertz law holds or whether there is indeed evidence that deceleration occurs. This section sets out some of the key contributions to this debate.

To illustrate that the idea of mortality deceleration is not new, Greenwood and Irwin (1939) suggests a high age plateau, albeit the relevance of their data to the current time is questionable. Vaupel (1979) proposes heterogeneous mortality levels between population groups as a cause of mortality deceleration and provides a theoretical justification for mortality deceleration based on Swedish data. Other studies which observe deceleration include Horiuchi and Wilmoth (1998), Thatcher (1999), Thatcher et al. (1998) and Wilmoth (1995). A number of these papers cite possible reasons for the observed deceleration, most of which focus on the heterogeneity of the data being analysed (see section 6.1.7 for an illustration of this effect on the mortality curve).

Whilst most authors acknowledge deceleration in mortality data and have offered up reasons as to what is driving this, two authors (Gavrilov and Gavrilova) have made a significant contribution to the debate with a robust defence that the Gompertz formula is appropriate for modelling μ_x (i.e. log μ_x is linear with age) at least to age 106 in homogeneous populations with clean data. We highlight here that 106 is lower than the limit age other authors have considered when looking at mortality deceleration.

We note that in their 1991 paper, these authors observe mortality deceleration at advanced ages, proposing that a logistic function might be a better fit for mortality at older ages. However, in more recent papers (2011 and 2014) they have put forward a clear view that deceleration is not observed up to relatively high ages. Their 2014 paper is a comprehensive summary of academic work in this area and itself looks at mortality trajectories based on the US Social Security Administration Death Master File (DMF) as well as data from the Human Mortality Database (HMD) for 1890 – 1899 birth cohorts. They also looked at the mortality curves of railroad workers versus the rest of the population since these are easily identifiable in the DMF.

Their 2011 paper looks at the impact on observed (and fitted) mortality rates for datasets with known differences in quality and shows deceleration is more prevalent in the lower quality data (see section 6.1 for worked examples of known data issues). There is also a helpful discussion regarding heterogeneity and the deviance in μ_x and q_x, especially at older ages.

In summary, Gavrilov and Gavrilova propose the following reasons why deceleration is observed in other studies (including their own), but that a Gompertz form is still the best fit for the underlying process:

- Age misreporting has been shown to lower the mortality curve at older ages (see section 6.1.5).

- Aggregation of single year birth cohorts – the authors have looked at the force of mortality for single year birth cohorts and demonstrate that if heterogeneity is introduced across cohorts, deceleration can be observed when we look at the combined experience of the birth cohorts.
- Studying age-specific probabilities of death rather than the force of mortality will lead to observations of deceleration at very high ages. Some papers do not provide clarity on the distinction between modelling in μ_x and q_x at old ages. Furthermore, they illustrate that studies using wide age intervals that assume a uniform force of mortality in each age interval produce bias downward to estimates of μ_x at higher ages.
- Older studies used data from birth cohorts where data recording was less accurate, whereas for more recent studies the data is more accurate.

The authors also note that there are other recent studies that support no deceleration, namely Stauffer (2002) and Bourbeau and Desjardins (2006).

Picking up on the idea of controlling for heterogeneity between birth cohorts, Beard (1971) looked at the theoretical form for μ_x for a heterogeneous population of subgroups where the force of mortality for a subgroup with "longevity factor" s is μ_x^s and $\varphi(s)ds$ is the proportion of the initial population with factor s. The force of mortality for the whole population can be written as:

$$\mu_x = \frac{\int \varphi(s)\mu_x^s \exp(-\int_0^x \mu_t^s dt)\, ds}{\int \varphi(s) \exp(-\int_0^x \mu_t^s dt)\, ds}$$

If we assume that each population has a Gompertz force of mortality (where s affects the level of $log\mu_x^s$):

$$\mu_x^s = \beta s exp(\lambda x)$$

and that $\varphi(s) = \kappa s^p \exp(-\gamma s)$ is a Gamma distribution, then it can be shown that the force of mortality for the population can be written as:

$$\mu_x = \frac{(p+1)\beta exp(\lambda x)}{(\gamma\lambda - \beta) + \beta exp(\lambda x)}$$

With a little rearrangement this can be written as a logistic form, which has been shown previously to allow for mortality deceleration at older ages. The above is, as Beard describes it to be, an interesting demonstration of how combining homogeneous groups can give rise to the logistic form (and hence mortality deceleration) for the aggregate population. However, he notes that "it suffers from the disadvantage that it assumes that the mortality of the various strata is Makeham in form for which experimental verification seems very difficult". (Note that although Beard refers to Makeham, we used the simpler Gompertz form in our example, above.)

In more recent work and in a similar vein, Currie (2010) suggests extrapolating mortality curves at high ages effectively using a different Gompertz curve for each year which leads to a decelerating curve when averaged through time. He also looks at fitting splines to μ_x at older ages which allow for more deceleration driven by the data at younger ages.

4.2. Mortality Deceleration – examples of Gompertz, Kannisto and Logistic extensions

To illustrate the practical impact of the above discussion, we set out below a comparison of Kannisto and Gompertz type extensions to the S2PML table. In each case the extension is applied from age 95, consistent with the methodology used in constructing that table.

Functional forms for Gompertz, Kannisto and Logistic type extensions

- **Gompertz extension.** The traditional Gompertz law assumes that the log of force of mortality is a linear function of age. Algebraically this can be expressed as $\mu_x = ae^{bx}$. This approach implies that the force of mortality continues to increase (without bound) as age increases.
- **Kannisto extension.** The Kannisto extension assumes that the logit of force of mortality is a linear function of age (where logit(μ_x) is defined as log(μ_x/(1-μ_x)). Algebraically this can be expressed as $\mu_x = ae^{bx}/(1+ ae^{bx})$. This approach results in the force of mortality being bounded, and converging towards 1 at very high ages.
- **Logistic extension.** We also consider the more general logistic functional form, which can be defined as $\mu_x = c + ae^{bx}/(1+ \alpha e^{bx})$. Like the Kannisto form, this approach results in the force of mortality being bounded, but in this case it converges towards $c + a/\alpha$ at very high ages.

Note that both Gompertz and Kannisto forms are special cases of the logistic form. The Gompertz form can be obtained by setting $c = \alpha = 0$; the Kannisto form is obtained by setting c = 0 and $\alpha = a$.

Regression versus extrapolation

For each of the above functional forms, we illustrate two methods of fitting the extension, which we have labelled as “regression” and “extrapolation” respectively.

Regression to determine high age extension

Under the regression approach, the high age extension is obtained by fitting the chosen functional form to the crude μ_x observed over the age range 80 to 95. The fitting approach is to minimise the weighted sum of square of errors. The error terms are taken to be the difference between the log (for Gompertz form) or logit (for Kannisto or Logistic forms) of the crude and fitted mortality rates. The weights used are the reciprocal of the variance of crude q_x.

Extrapolation to determine high age extension

With the extrapolation approach, the high age extension is determined using constraints at age 95 (i.e. where the extension meets the graduated rates) based on matching μ_x and derivatives of μ_x at that join. In the case of the Gompertz and Kannisto forms (which have two parameters) the constraints are to match μ_x and the first derivative of μ_x. The Logistic form has four parameters and so would require the matching of μ_x and the first three derivatives of μ_x which seems excessive, so in practice the “*c*” parameter has been set to zero and the constraints are to match μ_x and the first two derivatives of μ_x.

Figure 4.1: Comparison of q_xs at high ages under different extension methods

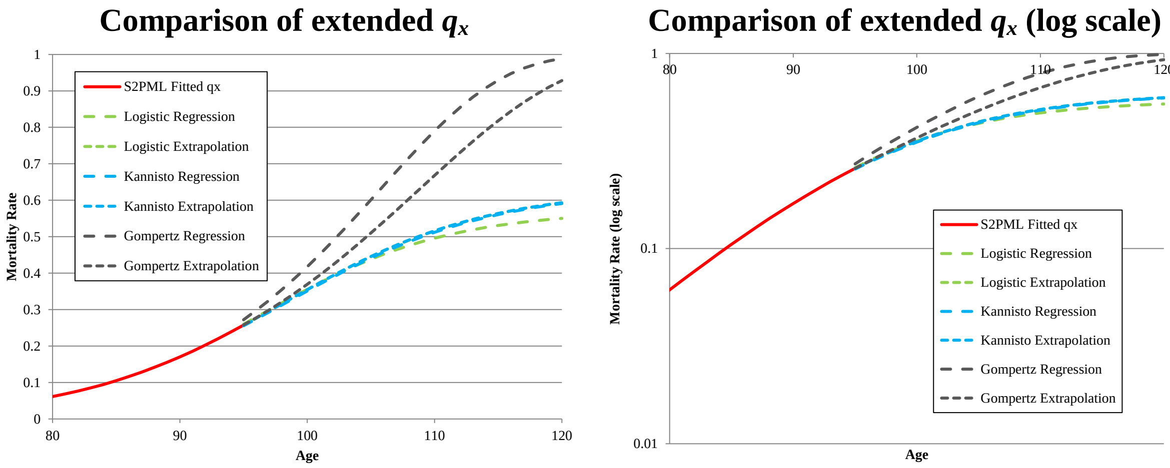

Figure 4.1 illustrates visually the extensions under each of the approaches described. Note that the Gompertz extensions are not a straight line on the log scale as the graph shows q_x values, whereas the Gompertz extension has been applied to μ_x.

Period life expectancies at sample ages under each extension are shown in Table 4.1. (Cohort life expectancies have also been considered but the results are similar, and the conclusions the same as those for period life expectancies, so these are not shown in the paper).

Table 4.1: Period life expectancies under different extension methods (and value relative to the "S2" Series extensions)

Period life expectancy	Logistic Regression	Logistic Extrapolation	Kannisto Regression	Kannisto Extrapolation	Gompertz Regression	Gompertz Extrapolation
e_{80}	8.027 (100.0%)	8.029 (100.0%)	8.032 (100.0%)	8.029 (100.0%)	8.001 (99.7%)	8.024 (99.9%)
e_{90}	4.036 (99.9%)	4.041 (100.0%)	4.050 (100.2%)	4.041 (100.0%)	3.958 (97.9%)	4.028 (99.7%)
e_{100}	2.125 (100.5%)	2.118 (100.1%)	2.142 (101.3%)	2.117 (100.1%)	1.698 (80.3%)	1.972 (93.2%)
e_{110}	1.486 (109.8%)	1.404 (103.7%)	1.417 (104.7%)	1.403 (103.7%)	0.753 (55.6%)	0.966 (71.4%)

The difference in high age values resulting from Gompertz versus Logistic/Kannisto extensions of S2PML are immediately apparent, with the Gompertz extensions resulting in clearly higher mortality rates (and lower life expectancies, particularly from age 100 and 110). This illustrates how the Logistic/Kannisto extensions decelerate relative to the Gompertz form, as described earlier in this section.

In this example, with the extensions applying from age 95, the differences in period life expectancies are relatively small up to age 90. However from age 100 the differences are marked with the Gompertz forms resulting in life expectancies around 10-20% lower than those under the Logistic/Kannisto extensions at age 100.

The Logistic and Kannisto extrapolations (and Kannisto regression) all result in similar mortality rates and life expectancies. The Logistic regression provides slightly lower mortality rates at the highest ages, and higher life expectancy from age 110.

It is noticeable in the above graphs and tables that there is a much greater difference between the regression and extrapolation extensions when using the Gompertz form than when using either Kannisto or Logistic forms. This appears to be driven by the Gompertz regression not matching the fitted S2PML q_x at the join (i.e. age 95) as well as the other extensions, including the Gompertz extrapolation, do. Examining the crude μ_x (and alternatively the fitted S2PML μ_x), there is evidence of "mortality deceleration" between ages 80 and 95. This means that while the Gompertz regression fits the crude μ_x closely over the age range 80-90 (where there is most data), it results in μ_x that are higher than the crude μ_x in the top part of that age range, particularly up towards age 95. This is why the Gompertz regression does not match the fitted S2PML at age 95 as well as the other extensions, and, as a result, why there is a greater difference between extensions based on Gompertz regression and Gompertz extrapolation.

4.3. Survey of other theories and studies

Mortality rates observed in other species

Mortality rates observed in mammals that exhibit similar survivorship patterns to humans can help inform the correct shape of mortality curve. The studies considered have looked at animals in captivity as this removes predators and controls for access to food and medical care. Species that experience high survival in early and middle life, followed by a rapid decline in survivorship in later life are considered due to their similarity with humans. The three papers reviewed conclude the following:

- Bronikowski et al (2011) conclude that, with a few exceptions, the Gompertz function is a good fit to the range of captive primates studied.
- Gavrilova & Gavrilov (2014) consider the shape of mortality for two species compared to 22 cohorts of humans, namely rats (10 cohorts) and mice (8 cohorts). They conclude that mortality deceleration at advanced ages is not a universal phenomenon, and survival of mammalian species follows the Gompertz law up to very old ages (106 for humans).
- Kohler et al (2006) present life tables by single year of age and sex for groups of animals and for 42 (mostly mammalian) species. At very high ages there is little information. The authors have used the Gompertz function to estimate high mortality and state briefly that they believe it is a good fit to the data available; we note that they did not fit any other functional forms to the data.

Broken Heart Syndrome

Various studies have shown that mortality rates for last survivors can be affected by the timing of the first life death. It is difficult to use this theory to help inform the general shape of the mortality curve at higher ages but it is noted as a potential contributing factor.

Broken heart syndrome (BHS) is seen as being a subset of BHS stress-induced cardiomyopathy or takotsubo cardiomyopathy. Typically a person is exposed to an extreme surprise or stressful event such as the loss of a loved one.

Carey et al (2013) demonstrate the statistically significant increased prevalence of myocardial infarction amongst those who have suffered the loss of a loved one. Carriere and Valdez (1996) illustrate that the pricing of joint life annuities can be reduced by up to 5% by allowing for dependant mortality models (compared with assuming independence).

The Working Party has not sought to develop this thinking further into making inference about setting mortality rates at old ages but wanted to highlight this as an area of potential relevant thinking on the subject.

Estimates of absolute mortality rates at very high ages

The estimates of mortality rates at very high ages later in this paper mostly reference the International Database of Longevity (IDL), which was an initiative set up in 2002 to combine validated international data regarding supercentenarians (those over age 110). This would appear to be the best dataset available to help inform mortality rates at very high ages.

Gampe (2010) is a study of mortality rates after age 110 using data from the IDL which covered 637 individuals, most of whom were US females. The results of the analysis were that mortality rates after 110 are flat with $\mu = 0.7$ implying a probability of death of $q_x = 0.5$. The authors found no sex specific difference or time trend observed in the studies performed.
This view is also supported by Kestenbaum and Ferguson (2010), who specifically look at 325 US supercentenarians born between 1870 and 1889. The validation of ages and dates of death of these lives is thorough, using a number of reliable sources. They looked at the probability of death above age 110 using an extinct generation methodology. Their results show a broadly constant probability of death of 0.5 for ages in advance of 110.

Another study by Robine, Gampe and Vaupel (2005) using the IDL data suggests that there is a constant force of mortality after age 110, leading to an annual death rate close to 0.5.

It is difficult to quantify the uncertainty around the estimates set out above since the data volumes underlying them are relatively small. However, it is reassuring that the three contributors above all have produced results that are consistent, giving readers some comfort that these parameters are somewhat robust.

The point should also be made that Gavrilova and Gavrilov do not make suggestions about the force of mortality for ages above 106 in their paper and the studies above are looking at those over 110, so although one purports no deceleration and the others propose a plateau, these do cover different areas of the curve.

5. Observed experience and graduated tables

5.1. Summary of graduated tables considered

To meet users' needs, mortality tables contain values to high ages – often to age 120, and almost always to well over 100. This requires tables to include some form of high age extension to cover ages where no experience data is available. In addition, at higher ages where some data *is* available, data volumes are frequently too low for robust statistical conclusions to be drawn and in some studies the data itself can be unreliable. Therefore, high age extensions also need to cover ages where there are low data volumes or the data is unreliable.

This section explores the data available, the (extended) graduated mortality rates, and the underlying extension methods at higher ages for a variety of UK and North American tables. We include both current and historic tables, as it is useful to observe the broad range of extension approaches that have been used now, and in the past.

The tables illustrated are described briefly below.

English Life Tables (ELT)

Based on population data from England and Wales combined, ELT17, ELT16, ELT15 and ELT14 reflect experience over the periods 2010-12, 2000-02, 1990-92 and 1980-82 respectively. No analysis of the method for graduating mortality rates at the oldest ages used for ELT17 is included in this working paper as ELT17 had not been published at the time of writing; however, Appendix B contains a brief description of the ELT17 methodology used to derive old age mortality rates.

CMI Self-Administered Pension Scheme (SAPS) tables ("S1" and "S2" Series)

These tables are based on data from UK occupational pension schemes. While a wide range of tables are provided, we focus below on the lives-based all-pensioner (excluding dependants) tables.

CMI Life Office tables ("80", "92", "00" and "08"Series)

These tables are based on data from UK life offices over the periods 1979-82 ("80" Series), 1991-94 ("92" Series), 1999-2002 ("00" Series) and 2007-2010 ("08" Series). Again a wide range of tables are provided; we focus below on the lives-based combined-pensioner tables.

Canadian Institute of Actuaries tables (CPM2014)

These tables are based on Canadian registered pension plan experience over the period 1999 to 2008, projected to 2014. The tables are amounts based.

Society of Actuaries tables (RP-2000 and RP-2014 series)

These tables are published by the Society of Actuaries' Retirement Plans Experience Committee and are based on uninsured private pension plans in the United States. The RP-2000 and RP-2014 tables reflect experience over the periods 1990-94 (projected to 2000) and 2004-08 (projected to 2014) respectively. The tables are primarily amounts based, although in the case of the RP-2014 tables, a headcount based variant is also produced and that has been used for the comparisons below.

Where both lives and amounts tables are available we focus on the lives variants; this provides consistency with the ELT tables which we use as a common comparator throughout. Key features are illustrated using the male data and tables; similar features apply to the corresponding female data and tables. Fuller details of the extension methods used in those tables are contained in Appendix B.

5.2. Graphical comparison of data underlying selected graduated tables

Figure 5.1 provides graphical comparison of the underlying data, the graduations and the extensions of the tables described in Appendix B. They also show the 2.5% and the 97.5% confidence limits (the "low gate" and the "high gate") of the crude values of q_x. ELT16 is provided as a comparator in each case.

Figure 5.1: Graduation and data underlying ELT16

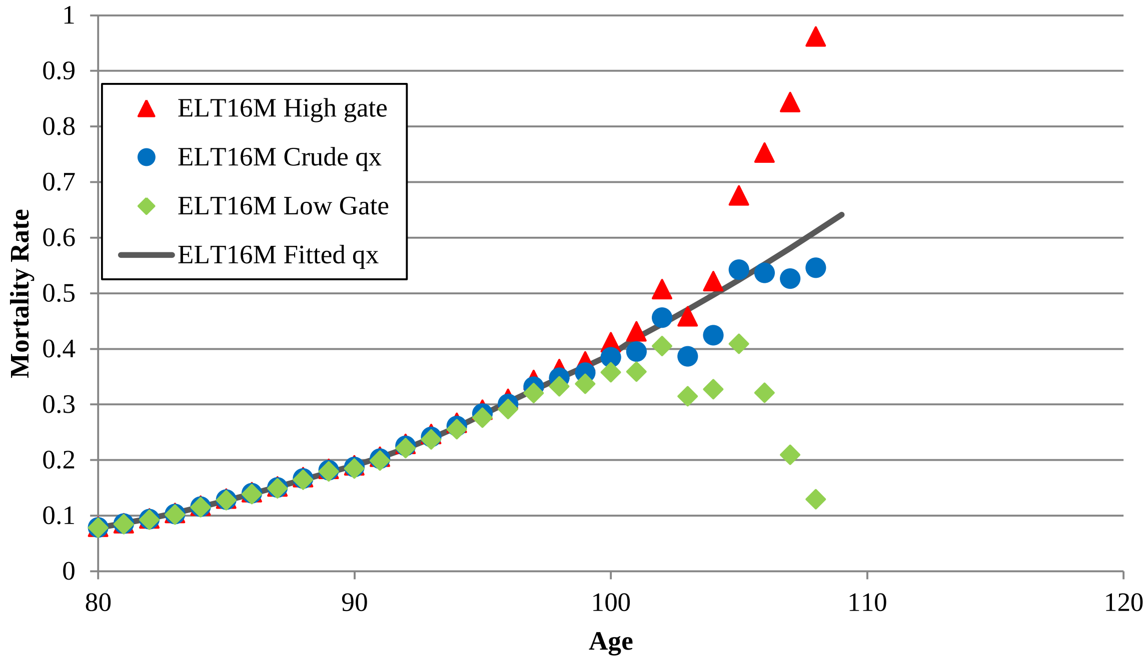

ELT16M uses data covering ages up to 108. Graduated mortality rates are determined using a variable-knot spline regression approach. There is no separate high age extension methodology; high age values are determined directly from the spline regression (with a high age constraint of $m_{120}=2$). With the exception of age 103 (where the crude q_x looks surprisingly low and the fitted rate exceeds the high gate) the fitted rates adhere reasonably closely to the crude data.

Figure 5.2: Graduation and data underlying S2PML

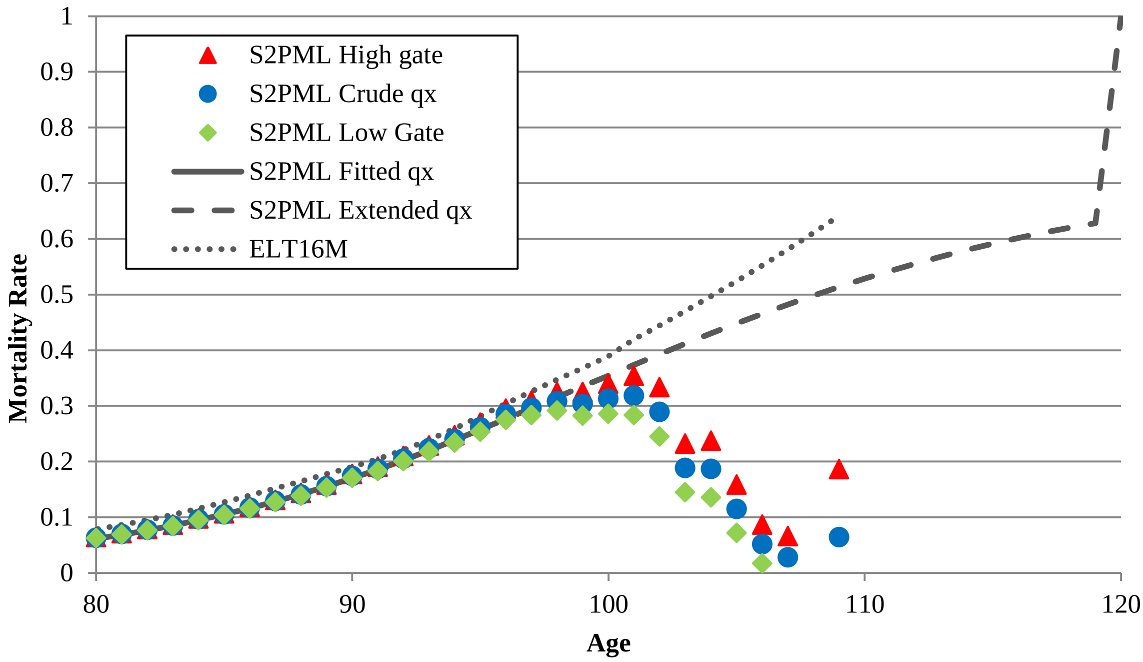

S2PML is graduated using data up to age 95. Figure 5.2 shows how the 95% confidence intervals around the crude q_x increase markedly above age 95 (reflecting the lower volumes of data at those ages). In addition the crude q_x decrease in value at ages around 100 and above (after levelling off in the high 90s), consistent with data distortions at higher ages (the potential causes of which are discussed further in section 6).

A high age extension provides q_x values at ages 96 to 120. This extension provides q_x values appreciably lower than ELT16M.

Figure 5.3: Graduation and data underlying PML08

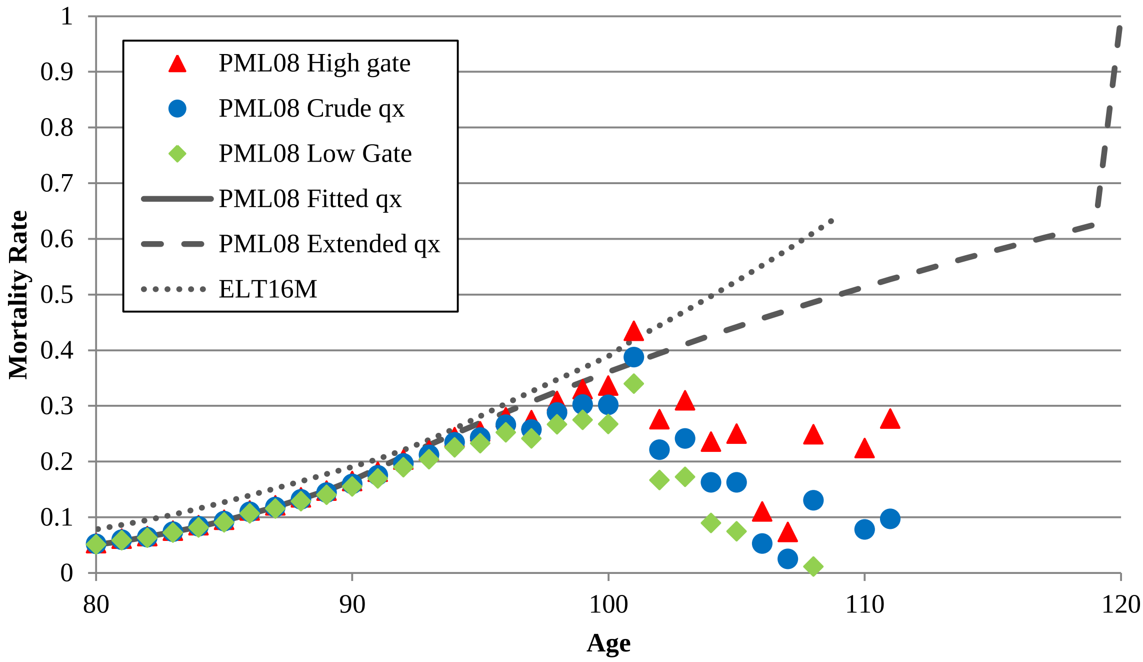

PML08 is graduated using data up to age 100. As for the SAPS data, Figure 5.3 shows how the 95% confidence intervals around the crude q_x increase markedly above age 95 (reflecting the lower volumes of data at those ages). The crude q_x decrease in value at ages around 100 and above (after levelling off in the high 90s); again consistent with data distortions at higher ages.

The high age extension provides q_x values at ages 90 to 120; these are appreciably lower than those for ELT16M.

5.3. Comparison of high age shape of graduated tables

Figures 5.4, 5.5, 5.6 and 5.7 provide a comparison of each family of tables, illustrating graphically both the graduated rates and the high age extension in each case. They demonstrate the approaches taken in practice and how these approaches have evolved over time.

Figure 5.4: Comparison of ELT tables

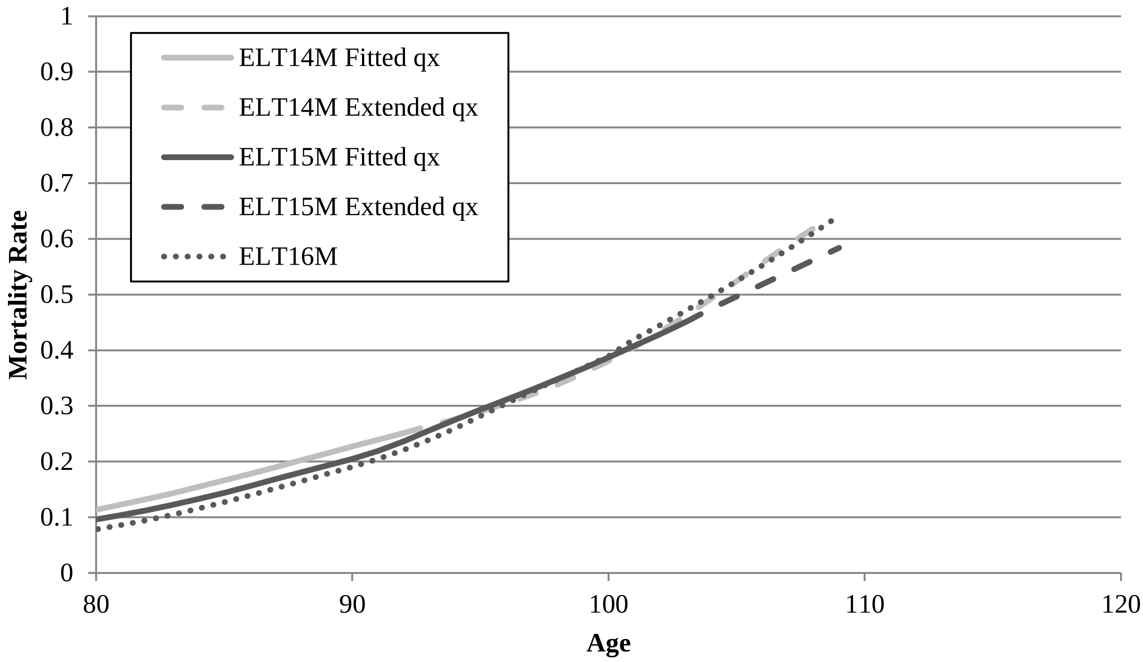

ELT14M and ELT16M contain similar high age values (despite relating to periods 20 years apart). ELT15M contains q_x values lower than ELT14M or ELT16M at high ages.

Figure 5.5: Comparison of CPM and RP tables

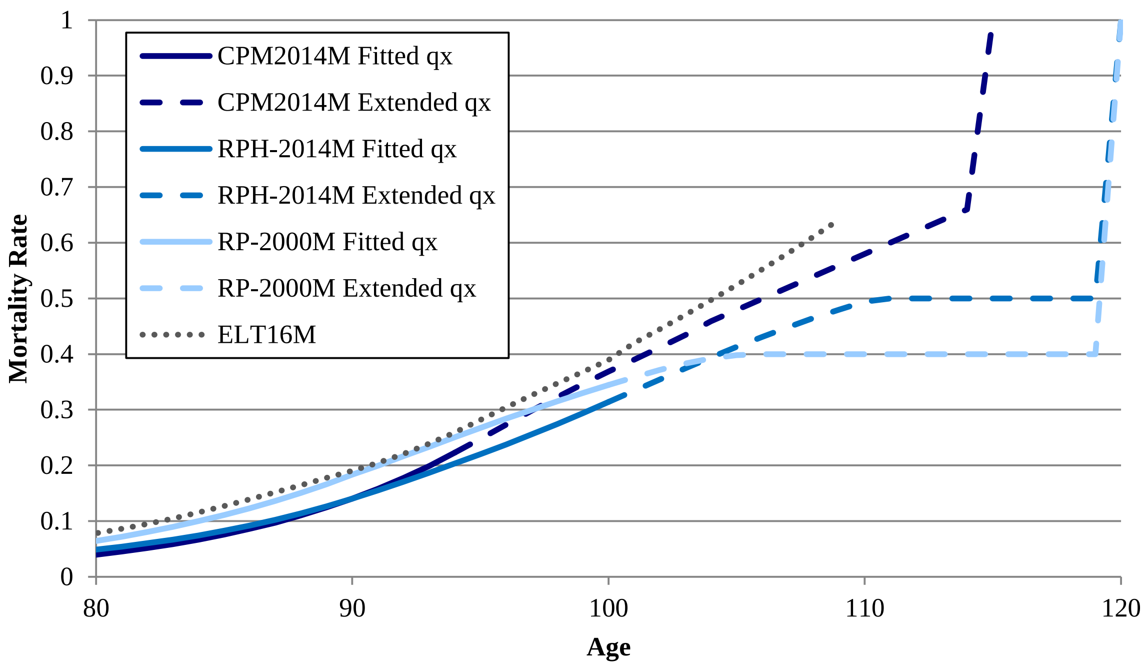

The high age q_x values are in each case lower than those for ELT16M. Those for RPH-2014M and RP-2000 are markedly lower, due to 0.5 and 0.4 limits being applied to q_x respectively.

Figure 5.6: Comparison of SAPS tables

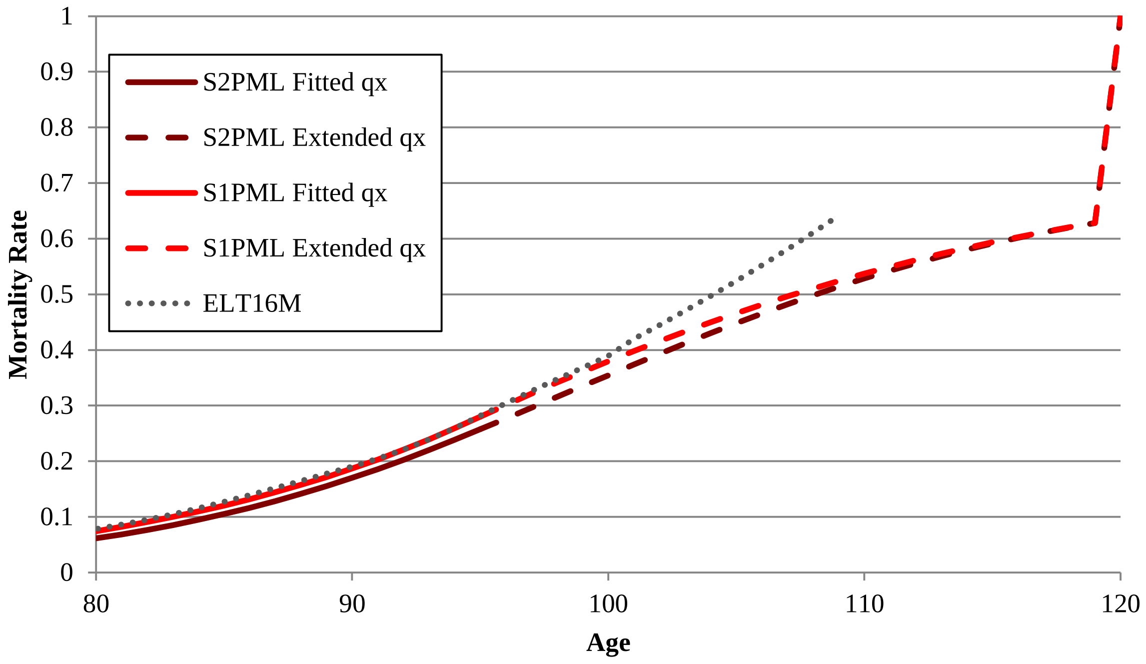

The “S1” and “S2” Series SAPS tables contain similar high age extensions. In each case the extension provides q_x values appreciably lower than and diverging from ELT16M.

Figure 5.7: Comparison of PML tables

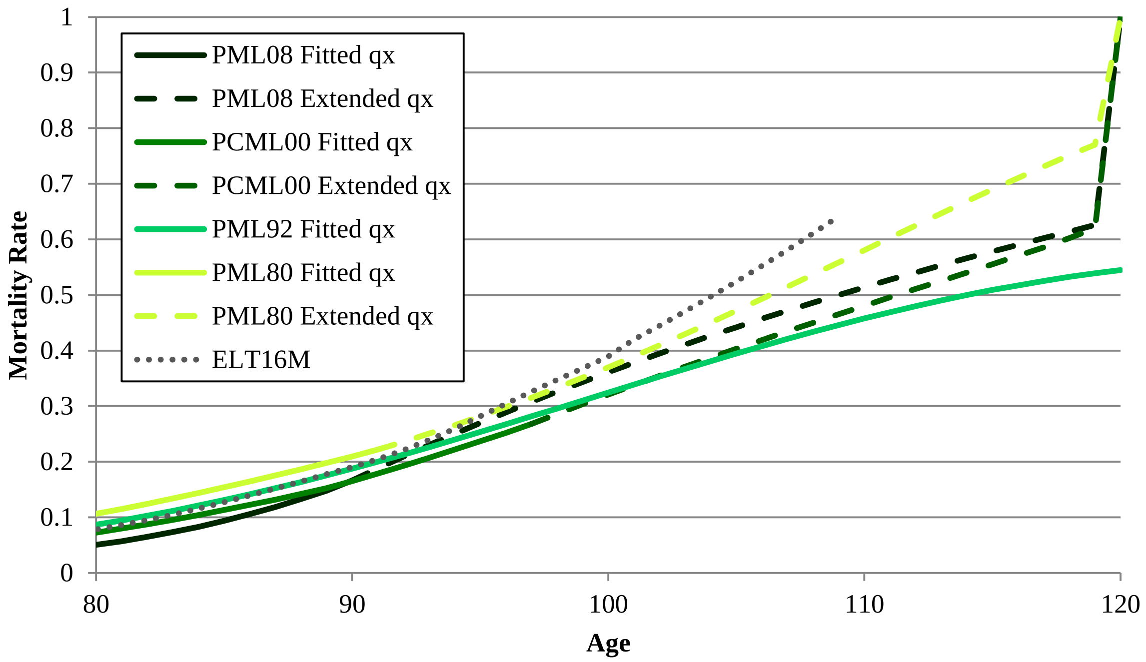

The high age q_x values are (particularly for PML08, PCML00 and PML92) appreciably lower than those for ELT16M. PML08 exhibits a markedly different shape to PCML00 and PML92 between ages 90 and 100; this reflects the intention for the early part of extension to take values close to those from UK population data.

5.4. Key features of different extension approaches

The table in Appendix B demonstrates the broad range of extension approaches adopted. This section explores briefly some of their key features.

Extension type

Continuation of graduation

The extension used for PML92 was essentially an extrapolation to higher ages of the fitted functional form from the main age range. This approach provides smoothness at the join between fitted and extended range, but considerable care (and potentially adjustment to parameters of the main fitting) is required to ensure values at the highest ages are reasonable.

Continuation of graduation (with constraints)

The extension of ELT16M was achieved by fitting the main graduation with an explicit constraint in the form of a high age limiting mortality rate. This is similar to the "continuation of graduation" approach but explicitly defines the mortality rate at the highest age.

Extension relative to another table or data source

PML80 was extended relative to another table (specifically PMA80). A similar concept was adopted for some of the "S2" Series tables. This approach can provide consistency between related mortality tables and can enable tables lower in a given hierarchy to inherit desired features from tables higher in that hierarchy. An independent extension methodology will however be required for at least one of the related tables, and this approach raises the impact of any inaccuracy in the core table. CPM2014 uses a variant of this approach, with a graduation of population data providing the upper part of the high age extension.

Extension using a functional form

Many of the tables (including PCML00, PML08, S1PML, S2PML, ELT14M and ELT15M) are extended using a specified functional form which applies only to the extended range. Functional forms used include cubic splines, logistic and others. Constraints are required to derive the parameters of the functional form (and are discussed further below).

Age from which extension applies

A broad range of ages are used for the age from which the extension applies (ranging from ages 90 to 103). This is typically set after consideration of the age above which data appears less reliable and / or data volumes reduce to an unsatisfactory level.

High age limiting values or constraints

Many of the extension methods include a high age limiting value.

- S1PML, S2PML and PCML00 adopted $\mu_{120} = 1$.
- RPH-2014M and RP-2000M adopted maximum q_x of 0.5 and 0.4 respectively.
- ELT16M adopted the rather higher limiting value of $m_{120} = 2$. This was obtained by setting $q_{120} = 1$ and using the approximation $q_x = m_x / (1 + m_x/2)$ which holds for $x=120$ if l_x is assumed linear in the age interval [120,121].

- ELT14M adopted what was described as "the somewhat arbitrary" value of $m_{105} = 0.75$ for males and 0.65 for females. This was derived from data collected for an investigation of the mortality of centenarians covering the years 1950 to 1979.
- Several of the tables are artificially closed at some age, e.g. q_x is set to equal 1 at age 112 in ELT14.

Other constraints

A range of other constraints have been adopted:

- μ_x (or some other mortality rate metric) is continuous where the fitted curve meets its extension. A constraint of this form is used in almost all the examples above (ELT15M and RPH-2014 being the exceptions, although the latter then used blending at the join) to avoid a step change in mortality rates.
- First differential of μ_x (or some other mortality rate metric) is continuous where the fitted curve meets its extension. This has the appeal of avoiding a step change in slope at the join. Care is required however as this constraint could serve to extrapolate edge effects.
- ELT14M effectively required the second differential of m_x to be continuous where the fitted curve met its extension. CPM2014 applies a similar form of constraint (by requiring q_x to be matched at 3 consecutive ages). This form of constraint provides additional smoothness at the join. However significant care is required as the second differential of any mortality rate metric at the edge of the fitting may have a very large parameter uncertainty, and so this approach may well extrapolate edge effects.
- First differential of μ_x (or some other mortality rate metric) is set to a specific value at the age associated with the limiting mortality rate. This may be a somewhat arbitrary assumption (i.e. will need to be set pragmatically rather than based on specific evidence), but can increase the degree of consistency between related tables.
- A "curvature" parameter, as used in the "00" and "08" Series extensions. This can provide a pragmatic means of adjusting the shape of the extension rather than being set based on specific evidence.
- Caps on q_x are a feature of both RPH-2014M and RP-2000M tables.

5.5. Comparison of the impact of different extension approaches

The impact of different extension approaches is illustrated in this section by applying these to the S2PML table. Seven alternative extension approaches have been considered:

- "S2" Series extensions (as used in the published S2PML table);
- "08" Series extensions. In particular those for PML08 (with a run-in age of 90 and curvature parameter of 0.95);
- "00" Series extensions, focusing on PCML00 (with a run-in age of 97 and curvature parameter of 0.8);
- RP-2014 extension approach from age 100, using a Kannisto regression (of crude μ_x) fitted over age range 75 to 104 and a maximum q_x of 0.5;
- RP-2000 extension approach from age 100, using a cubic polynomial and a maximum q_x of 0.4;

- ELT15M extension approach (using logistic regression of fitted rates for ages 85 to 103 to determine extension); and
- ELT14M extension approach (with extension from age 92, as per ELT14M).

The results of these extensions are shown graphically in Figure 5.8, along with sample period life expectancies, in Table 5.1.

Figure 5.8: Comparison of extensions of graduated tables at high ages

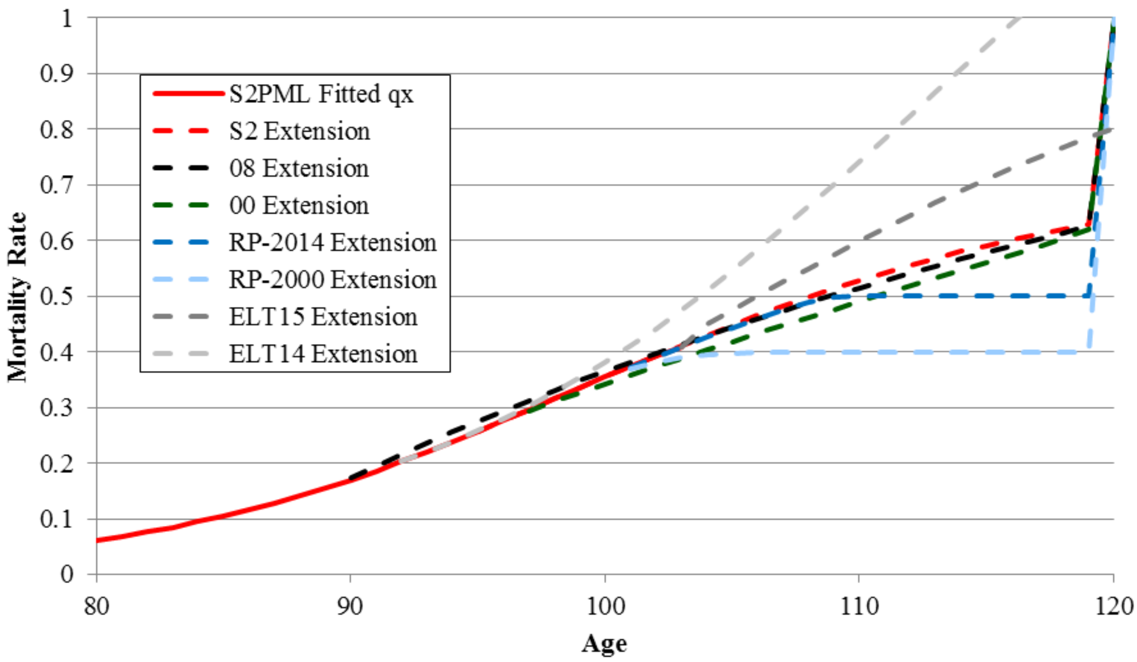

Table 5.1: Period life expectancies under extensions of different tables (and value relative to "S2" Series extensions)

Period life expectancy	"S2" Series	"08" Series	"00" Series	RP-2014	RP-2000	ELT15	ELT14
e_{80}	8.029 (100.0%)	7.973 (99.3%)	8.032 (100.0%)	8.029 (100.0%)	8.030 (100.0%)	8.029 (100.0%)	8.020 (99.9%)
e_{90}	4.041 (100.0%)	3.874 (95.8%)	4.052 (100.3%)	4.042 (100.0%)	4.045 (100.1%)	4.040 (100.0%)	4.014 (99.3%)
e_{100}	2.115 (100.0%)	2.076 (98.1%)	2.227 (105.3%)	2.129 (100.7%)	2.192 (103.6%)	2.101 (99.3%)	1.888 (89.3%)
e_{110}	1.354 (100.0%)	1.399 (103.3%)	1.487 (109.8%)	1.499 (110.7%)	1.991 (147.1%)	1.133 (83.7%)	0.825 (60.9%)

We discuss below the impacts of the alternative extension approaches.

"S2" Series Extension

This is as was adopted for the "S2" Series tables.

"08" Series Extension

Some of the high level features of this extension are similar to those from the "S2" Series extension. In particular, the same high age limiting value is assumed.

However, this extension gives by far the lowest life expectancy from age 90, some 4% lower than that under the "S2" Series extension. This is driven by the extension applying from a younger age than other extensions (i.e. from age 90), combined with a different curvature to other extensions, which results in the mortality rates from age 90 to age 100 being higher than under other extension approaches. As noted previously, the parameters were selected to result in values close to those from UK population data in the early part of the extension. In addition, the way the "08" Series rates approached population mortality rates was considered when the "08" Series parameters were set, hence had that approach been applied to the "S2" Series tables then different parameters may have been adopted.

"00" Series Extension

The high level features of this extension are similar to those from the "S2" Series extension. In particular, the same high age limiting value is assumed. The extension method is closer to a straight line (in terms of q_x versus age x) than that under the "S2" Series approach, hence mortality rates are slightly lower and life expectancies slightly higher.

RP-2014 Extension

Graphically, this extension appears significantly different to most others due to the maximum q_x value of 0.5. However the fact that this extension only applies from a relatively high age (i.e. 100) and the limiting value (of 0.5) is not attained until around age 110 results in life expectancies at ages 80, 90 and 100 that are very similar to those under the "S2" Series approach.

RP-2000 Extension

The most striking feature of this extension is the maximum q_x value of 0.4. This results in:

- a higher life expectancy from age 100 than most other approaches (although lower than that under the "00" Series approach); and
- by far the highest life expectancy from age 110 (nearly 50% higher than that under the "S2" Series approach).

ELT15 Extension

The extension applies from a higher age (i.e. 103) than the other extensions. As the "S2" Series extension is in fact very similar to the underlying "S2" Series fitting to that age, the ELT15 approach results in very similar mortality rates up to age 103.

It can be seen from Figure 5.8 that there is a step change in q_x at age 103-104, at which point the extension starts to apply. This is a result of the regression approach not requiring the extension to have the same value as the underlying fitting at the join.

The high age values are driven by extrapolation of the logistic regression (rather than a high age constraint assumption). In this case this results in high age mortality rates clearly higher than the “S2” or “00” Series extensions.

Despite the high age differences, sample life expectancies are similar to those under the “S2” Series extension approach at ages 80, 90 and 100.

ELT14 Extension

As for the ELT15 extension, the very high age values are driven by an extrapolation of the fitted table (plus a constraint at age 105), with no constraint on the value at the limiting age. In this case this approach results in high age mortality rates clearly higher than the “S2”, “00” Series or ELT15M extensions.

Despite the very clear high age differences, sample life expectancies are similar to the other approaches at ages 80 and 90, with appreciable differences occurring only from age 100. Cohort life expectancies have also been considered. The results are similar and the conclusions the same as those for period life expectancies.

6. Analysis of high age mortality

6.1. Illustration of common issues in data and modelling

6.1.1. Estimation of exposure

In order to calculate mortality rates at any age, data are needed on the numbers of deaths and on the exposed to risk. Depending on the model adopted either a central or an initial exposed to risk may be derived from the data. There may be additional factors to consider when calculating exposed to risk for the oldest ages which may not be significant at younger ages. These can vary depending on whether this is being calculated using life office data or pension scheme data for a single entity or aggregated data or population data; there may also be different issues to be considered in the deaths data available for each of these.

Life office and pension scheme data

For most actuarial investigations related to life office or pension scheme mortality experience exposure would be calculated at each age. Where the date of birth, date of entry (to the exposure being calculated) and date of death (or other movements of individuals in the experience) are known, it is usually possible to calculate the amount of time individuals have been exposed to the risk of death during the investigation period at each age to the nearest day. These data can then be aggregated to give exposed to risk by age over the period under consideration.

In theory it is possible to track policyholders or pension scheme members over time whilst they are alive and record their eventual dates of death or other exit from the database. In particular, the deaths data should be linked to the exposed to risk and it is unlikely that a death would be recorded in respect of a person who had not been recorded as an in force case at some earlier date. The datasets available should also give a complete count of all the individuals covered, albeit there may be duplicates.

Population data

For population data at a national level, exposed to risk is generally derived using a census method based on population estimates by single year of age, since data at an individual level are not generally available, unless the country operates a population register system.

Deaths data are usually provided on a calendar year basis and hence an exposed to risk by calendar year is also required. If population estimates are available at the mid-point of a calendar year, as for the UK, these estimates are often taken as the exposed to risk. If estimates are available at the beginning and end of a calendar year then these may be interpolated to obtain an estimated exposed to risk for that year. Using these assumptions effectively assumes a uniform birth rate profile or at least one in which the weighted average birthday is the mid-year point.

Calculating exposed to risk using population data provides, to some extent, a different set of problems from those encountered when calculating exposed to risk from life office or pension scheme data, at least as far as the UK is concerned.

Issues related to deaths data, calculating exposed to risk and mortality rates at the population level are discussed in more detail in section 6.2.

6.1.2. Derivation of mortality rates

In order to derive mortality rates, deaths data is required in addition to exposed to risk. Mortality rates can be expressed as the force of mortality (or hazard rate) (μ_x), central death rates (m_x) or probabilities of mortality/initial death rates (q_x).

Often, central death rates are calculated directly from the data by dividing deaths at age x by the exposed to risk at age x, or the force of mortality may be derived by fitting a formula to the data. There are various approximations commonly used to convert between μ_x, q_x and m_x values, for example $m_x \approx \mu_{x+½}$, $\mu_{x+½} \approx$ -ln$(1 - q_x)$ (or alternatively $q_x \approx 1\text{-}e^{-\mu(x+½)}$) and $q_x \approx m_x/(1+m_x/2)$. These approximations may not be appropriate at the oldest ages due to underlying assumptions such as uniform distribution of deaths over the age interval, which (if not true) may cause material distortion in these approximations at high ages. The accuracy of all these approximations depends on the underlying assumptions about the force of mortality and how it changes over age.

Initial rates of mortality (q_x) may not be the most appropriate rates to use when considering mortality at the oldest ages, especially where these are calculated on annual intervals. Initial rates have an upper bound of 1.0. The application of the q_x rate may be correct in deriving the numbers of deaths for a given number at the beginning of the year of age, but it does not give any indication of the periods spent alive for those who die during the year – at the oldest ages deaths may not occur in a linear fashion between ages x and x+1.

6.1.3. Data Issues

There may be errors in the database, whether life office, pension scheme or population data, for an individual arising from:

- Missing date of birth;
- Misreporting of date of birth;
- Misstatement of age (if used instead of calculating from date of birth);
- Missing date of death;
- Misreporting of date of death;
- Delay in reporting death; and
- Non-reporting of death.

In addition, for life office and pension scheme data there may be errors due to deaths being assumed in some cases; for example, if bank accounts appear to have become inactive, or no contact has been made for several years so the person is assumed to be dead.

Errors in data for in force members will affect the calculation of the exposed to risk. Errors in the deaths data will affect the calculation of mortality rates through both the deaths data and the exposed to risk.

For cases where a date of birth is not known, a date may have been imputed or an arbitrary date may have been assigned. If a death is notified there should be sufficient detail to link the death to an in force record in the dataset, even if the reporter does not know the person's details.

Late or non-reporting of deaths may occur if it is not known that the person who died had a life insurance policy or a pension. At the oldest ages a significant overstatement of exposed to risk can occur even if deaths are reported on a timely basis if there are long delays in removing them from the 'in force' data.

Errors may also occur if data has been transferred from another life office or pension scheme or from administration errors.

The effects of age misstatement, which can in turn be linked to errors in dates of birth or death, and of late or non-reporting, have been modelled and are discussed later in the paper.

Any of the above errors are likely to have a greater effect on the calculation of exposed to risk at the oldest ages than at younger ages, as there are likely to be comparably fewer people at those ages in the database. Hence, any error is a greater proportion of those at the oldest ages than at younger ages. For instance if a death at age x is not reported (or for some other reason that person is not removed from the 'in force' database) then that person will effectively add a year of exposed to risk at each age as they age through time. These 'phantom' people will have a greater proportionate effect on the exposed to risk at older ages as there are fewer actual people alive at those ages; the measured exposed to risk will be higher than it should be and this will reduce the mortality rates calculated at these ages.

There may also be issues arising in respect of individuals residing overseas. This would not affect the calculation of the exposed to risk or mortality rates unless any deaths were not notified or there were longer delays before notification; in theory, deaths overseas should still be linked to an in force person.

Data issues such as those discussed above can be one reason why the rate of increase in mortality rates calculated from datasets, such as the CMI life office data, often appears to reduce or even turn negative at the oldest ages. The possible effects of delays in reporting deaths and in calculating ages are discussed in more detail in the next sections of this working paper.

Some of the above issues also apply when estimating exposed to risk and mortality rates at a population level; these are discussed in more detail in section 6.2.

6.1.4. Delays in death reporting

Delays in death reporting have a distorting effect on crude mortality rates for two reasons, the first being the understatement of deaths in the numerator and the second being the erroneous inclusion in the denominator of exposure from unreported deaths. At older ages, given proportionately lower exposures and higher death rates, these effects can be material.

As an illustration of this point, we present two scenarios in which we have solved for the levels of unreported deaths which would reconcile the crude rates observed in the CMI 2007-10 pension annuities in payment dataset (as used for the "08" Series tables) to underlying rates given by:

a) Mortality rates assuming a simple Gompertz law holds (log μ_x is linear in x); and
b) The mortality rates underlying the graduated rates at older ages (i.e. population data).

We have considered a simplified model tied to the crude rates underlying the PMA08 mortality tables as published in CMI Working Paper 81 with the following assumptions:

- Modelling period is 2002 to 2012.
- We have assumed a 'stationary' population, say 1,000 lives at age 65 in each year, reducing by age according to the assumed underlying mortality rates.
- Allowance for mortality improvements in line with the CMI_2014 model from 2001 onwards (no allowance for mortality improvements in earlier years in constructing the initial 'stationary' population).
- A proportion of deaths are assumed unreported by the observation date (end 2011).

Given the table is based on experience over 2007 to 2010 and uses deaths reported by 31 December 2011, we have assumed various delay patterns in death reporting for the period 2002 to 2011. This assumes that the population is known and defined in 2002 and so there are no long term unreported deaths prior to this year. We have then looked at what percentage of deaths would need to be unreported to observe the (known) crude rates if the underlying "true experience" were i) a Gompertz law or ii) were the graduated rates, as well as showing the sensitivity by illustrating the effect of other delay patterns.

These two illustrations are based on the observed experience by amounts, but the results are broadly similar (perhaps surprisingly) on a lives basis.

Distortion from Gompertz

Figure 6.1 illustrates graphically the distortion in shape of a Gompertz curve under several death reporting delay scenarios. In order to illustrate this point, we have fitted a simple Gompertz curve to log μ_x rates from the crude m_x rates for ages 65 to 85. The curve has been fitted by minimising the unweighted sum of squares of errors in log μ_x and using the approximation $\mu_{x+0.5} \approx m_x$. We have modelled assumed true experience on these assumed mortality rates (shown by the blue line) and then calculated the observed deaths and exposures allowing for 3.5% of deaths (from 2002 onwards) to be unreported as at 31 December 2011 (shown as mortality rates by the red line). The resulting distortion of observed experience is seen as the difference between the blue and red lines. The inference here, therefore, is that if the Gompertz law held for this group, and the only distortion were from late reporting of deaths, one would need 3.5% of deaths to be unreported (since 2002) to support the observed q_x rates (shown by the red line reasonably matching the purple crude experience 'data points' line).

Figure 6.1: Comparison of unreported death scenarios with Gompertz fit and crude rates

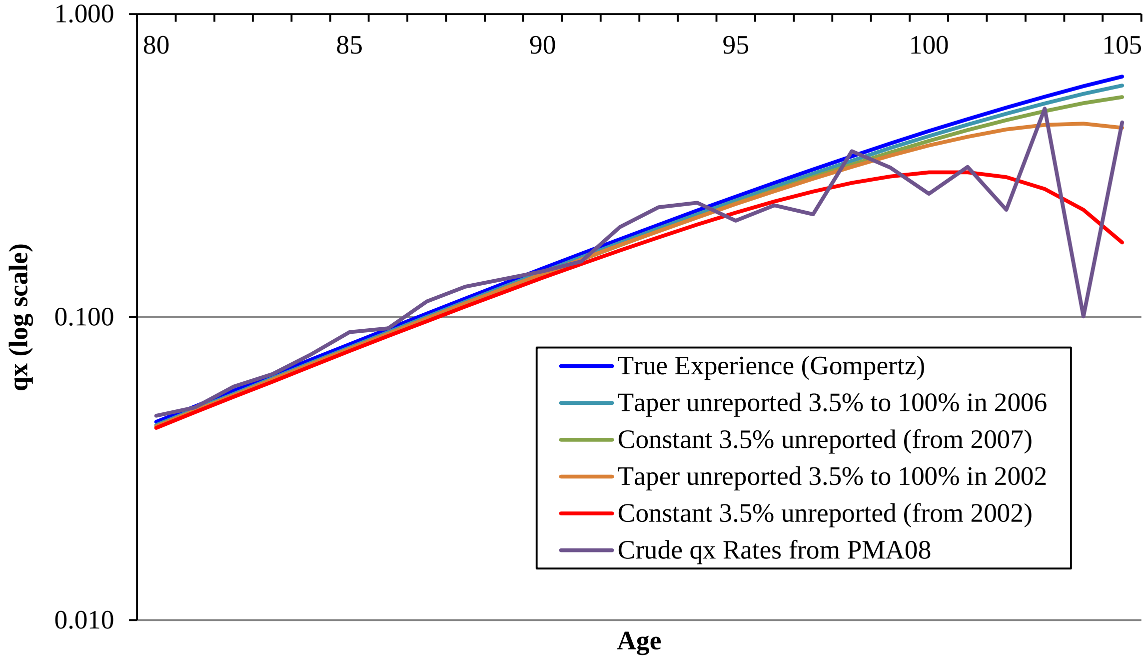

We have considered other less extreme scenarios to demonstrate the sensitivity to the levels of death reporting observed over time. Details of these are shown in Table 6.1 and on the graph as well as Table 6.2 which shows life expectancies under the various scenarios.

Table 6.1: Proportion of deaths reported by year under scenarios considered (Gompertz)

Scenario / Year	**2002**	**2003**	**2004**	**2005**	**2006**	**2007**	**2008**	**2009**	**2010**
Constant 3.5% unreported (from 2002)	96.5%	96.5%	96.5%	96.5%	96.5%	96.5%	96.5%	96.5%	96.5%
Constant 3.5% unreported (from 2007)	100.0%	100.0%	100.0%	100.0%	100.0%	96.5%	96.5%	96.5%	96.5%
Taper unreported 3.5% to 100% in 2006	100.0%	100.0%	100.0%	100.0%	100.0%	99.1%	98.3%	97.4%	96.5%
Taper unreported 3.5% to 100% in 2002	100.0%	99.6%	99.1%	98.7%	98.3%	97.8%	97.4%	96.9%	96.5%

Table 6.2: Gompertz life expectancies versus levels of late reporting

	Gompertz Curve	Constant 3.5% unreported (from 2002)	Constant 3.5% unreported (from 2007)	Taper unreported 3.5% to 100% in 2006	Taper unreported 3.5% to 100% in 2002
Period Life Expectancy					
65	20.33	20.76	20.63	20.51	20.60
80	9.17	9.56	9.39	9.30	9.38
90	4.31	4.75	4.49	4.41	4.50
Percentage difference versus Gompertz					
65	100.0%	102.1%	101.4%	100.9%	101.3%
80	100.0%	104.3%	102.5%	101.4%	102.4%
90	100.0%	110.1%	104.0%	102.2%	104.4%

Distortion from Graduated Rates

The second illustration is similar, but the "true experience" has been assumed to be in line with the final PMA08 q_x rates, which have been calibrated at older ages to trend towards population mortality rates. These rates are higher than the crude rates but decelerate on a log scale as can be seen in Figure 6.2. We refer to these rates in this paper as the "graduated rates". For this example we find an assumed non-reporting rate of 2% reconciles the underlying mortality rates to the actual observed crude rate. In line with the previous illustration, we also show two further sensitivities. Given the assumed level of misreporting is lower, we have had to extend the unreported deaths back to 2004 to move the curve materially down, and tapering from any year later than 2002 also does not move the curve materially.

Figure 6.2: Comparison of unreported death scenarios with Graduated and crude rates

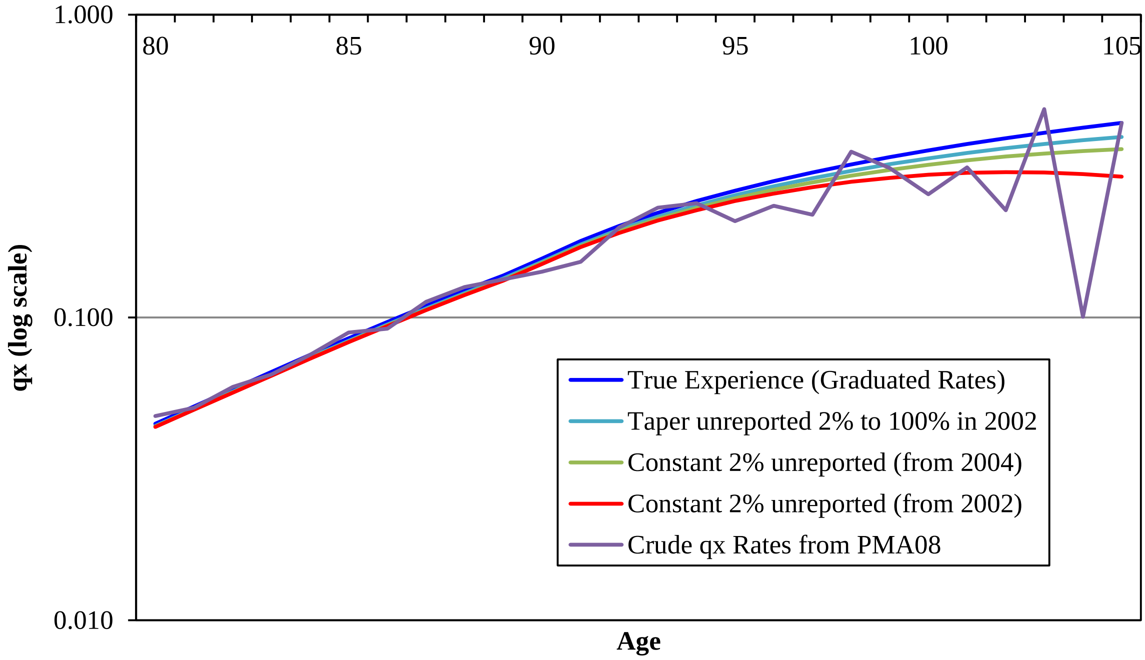

Table 6.3: Proportion of deaths reported by year under scenarios considered (Graduated)

Scenario / Year	2002	2003	2004	2005	2006	2007	2008	2009	2010
Constant 2% unreported (from 2002)	98.0%	98.0%	98.0%	98.0%	98.0%	98.0%	98.0%	98.0%	98.0%
Constant 2% unreported (from 2004)	100.0%	100.0%	98.0%	98.0%	98.0%	98.0%	98.0%	98.0%	98.0%
Taper unreported 2% to 100% in 2002	100.0%	99.8%	99.5%	99.3%	99.0%	98.8%	98.5%	98.3%	98.0%
Taper unreported 2% to 100% in 2003	100.0%	100.0%	99.7%	99.4%	99.1%	98.9%	98.6%	98.3%	98.0%

Table 6.4: Graduated Rates Life Expectancies versus levels of late reporting

	Graduated Rates	Constant 2% unreported (from 2002)	Constant 2% unreported (from 2004)	Taper unreported 2% to 100% in 2002
Period Life Expectancy				
65	20.22	20.46	20.43	20.37
80	8.95	9.17	9.13	9.07
90	4.10	4.36	4.29	4.22
Percentage difference versus Graduated Rates				
65	100.0%	101.2%	101.0%	100.7%
80	100.0%	102.5%	102.0%	101.4%
90	100.0%	106.3%	104.5%	102.7%

The analysis in Tables 6.3 and 6.4 illustrates that unreported deaths could have a significant distorting impact on the q_x curve at older ages. Of course, the graduated rates would also need to be 'grossed up' at younger ages for the assumed non-reporting but the effect is less marked.

The level of 3.5% required to distort an assumed Gompertz curve to the crude rates observed at high ages may be considered a little high, and the level of 2% needed to reconcile the graduated rates to the crude observations seems a little more reasonable. We note that an initial analysis of the Annuitant experience dataset by the CMI Annuities Committee suggested that around 0.9% of deaths are reported 3-4 years after the event and 0.6% 4-5 years after, so a rate of late reporting of 2% for years 4+ is not implausible (although in our central case we assume this level persists).

The analysis also shows that these effects can have of the order of 1-2% impact on life expectancy at age 65 depending on the level of distortion assumed, with the impacts being higher at older ages.

6.1.5. Misstatement of age

The purpose of this section is to investigate the impact on the mortality curve of misstatement of age in deaths and exposures data. Preston, Elo and Stewart demonstrate in their 1997 paper that regardless of net over-, symmetric- or understatement of age, misstatement tends to underestimate mortality rates at older ages, due to the slope of the age distribution at higher ages. This is because erroneous data flowing into older ages will have a larger effect than when flowing into younger ages (even if they are proportionately larger) due to the dramatic reduction in lives in force as age increases. In the examples below we consider symmetric misstatement of age only.

We have looked at the impact of age misstatement on mortality rates using a similar approach to that above, repeating both the Gompertz and graduated rates examples. We have looked at four scenarios where we allow for a given percentage of the observed deaths and exposures to have a misstated age that is uniformly distributed across a predefined age range relative to the age in question. We have looked at this analysis across ages 65 to 105 and for ages at the edge of the analysis, we have allowed for misstatement to be restricted to the ages in scope of the investigation (e.g. for age 67 with a 1% misstatement allowance +/- 10 years, we have re-distributed 1% of the deaths and exposures to ages 65-66 and 68-77 from those at age 67).

The analysis below shows the distortion to the assumed "true" experience for each of the examples. As for late reporting, we have included the impact of the adjustments on life expectancies.

Scenarios Considered

To help inform what a reasonable level of age misstatement might be, we have conducted some straw man analyses of date of birth (DOB) corrections resulting from data cleanse exercises on a number of large buy-in clients of a bulk annuity provider and due diligence findings from checks of dates of birth on system data against those recorded on death certificates, with the results shown in Tables 6.5 and 6.6.

Table 6.5: Date of birth changes as a result of data cleanses for 3 schemes

Data Cleanse Changes to in force DOB	Scheme 1	Scheme 2	Scheme 3
Number of members in scheme	25,693	5,813	8,811
Number of affected members	16	9	13
Percentage Affected	0.06%	0.16%	0.15%
Average Age of affected members	81	72	80
Average Age of Scheme	75	71	72
Correct DOB number of years younger (average)	0.66	2.33	2.29
Correct DOB number of years younger (std dev)	3.75	15.17	4.54

Table 6.6: Date of birth errors versus death certificates from due diligence exercises

Death Certificate checks of random sample of deaths	Insured Block 1	Scheme 4	Scheme 5
Number of records in sample	109	248	85
Number of affected members	3	2	1
Percentage Affected	2.75%	0.81%	1.18%
Average Age of affected members	83	87	105
Average Age of Sample	82	89	86
Correct DOB number of years younger (average)	-5	0.5	No DOB in pricing data
Correct DOB number of years younger (std dev)	7.88	0.71	n/a

The levels implied by the data cleanse analysis may understate the occurrence of age misstatement since changes to dates of birth are only made as the errors become evident as part of business as usual processes. The due diligence results, although much smaller samples, indicate a level of misstatement of something closer to 1%, with it being difficult to draw a strong conclusion on the spread of age misstatement given the low number of samples. In light of these results, we have included scenarios of 1% misstatement with a 10 and 20 year spread as scenarios at the worse end of the spectrum.

In Figures 6.3 and 6.4 and Tables 6.7 and 6.8, we have included a scenario of 8% age misstatement with a spread of +/-2 years as a representation of the level of age misstatement observed in Hill, Preston and Rosenwaike (2000) who completed a comprehensive study comparing ages on death certificates against social security records in the US. Although this is not directly comparable to the UK situation, with dates of birth being recorded on death certificates, it is helpful to see this level of "real world" misstatement alongside the other scenarios considered.

We have also included a scenario based on analysis of raw data provided to the Club Vita dataset (i.e. reflecting the age misstatements typically identified by Club Vita's routine data cleansing). From this we believe a reasonable assumption could be that 0.1% of data contains age misstatement which is distributed uniformly across the age spectrum (in our example this is ages 65 to 105).

Figure 6.3: Comparison of age misstatement scenarios with Gompertz fit and crude rates

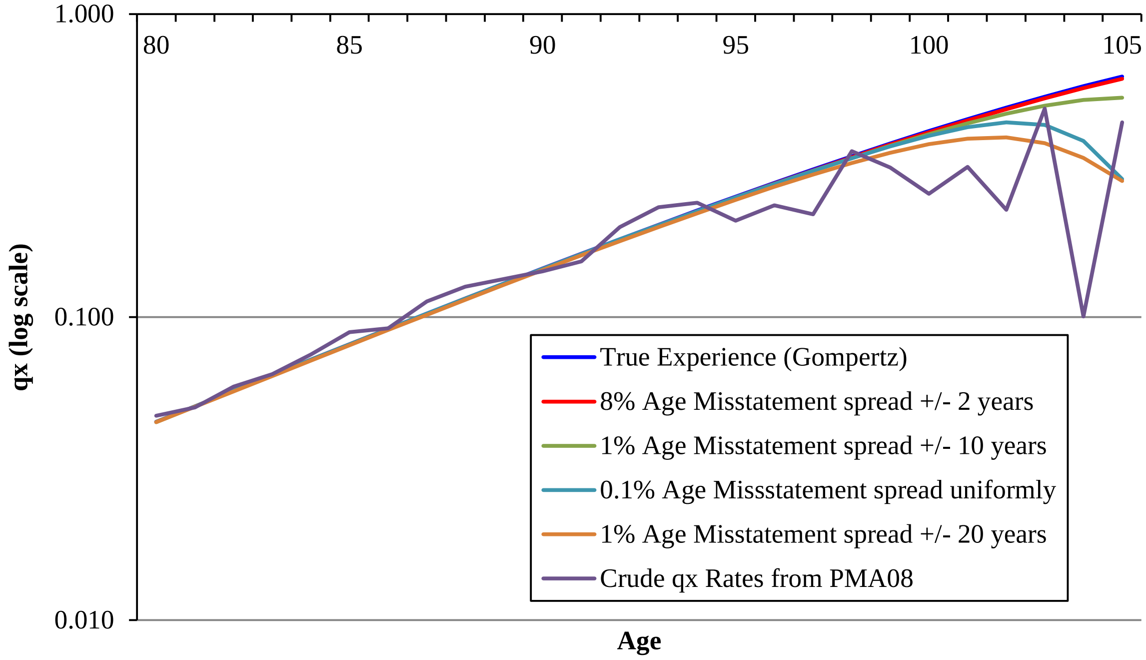

Table 6.7: Gompertz Life Expectancies versus levels of age misstatement

	Gompertz Curve	8% Age Misstatement spread +/- 2 years	1% Age Misstatement spread +/- 10 years	0.1% Age Misstatement spread uniformly	1% Age Misstatement spread +/- 20 years
Period Life Expectancy					
65	20.33	20.34	20.34	20.34	20.36
80	9.17	9.18	9.18	9.18	9.23
90	4.31	4.33	4.34	4.34	4.40
Percentage difference versus Gompertz					
65	100.0%	100.0%	100.0%	100.0%	100.1%
80	100.0%	100.1%	100.2%	100.1%	100.7%
90	100.0%	100.3%	100.6%	100.5%	102.1%

Figure 6.4: Comparison of unreported death scenarios with Graduated and crude rates

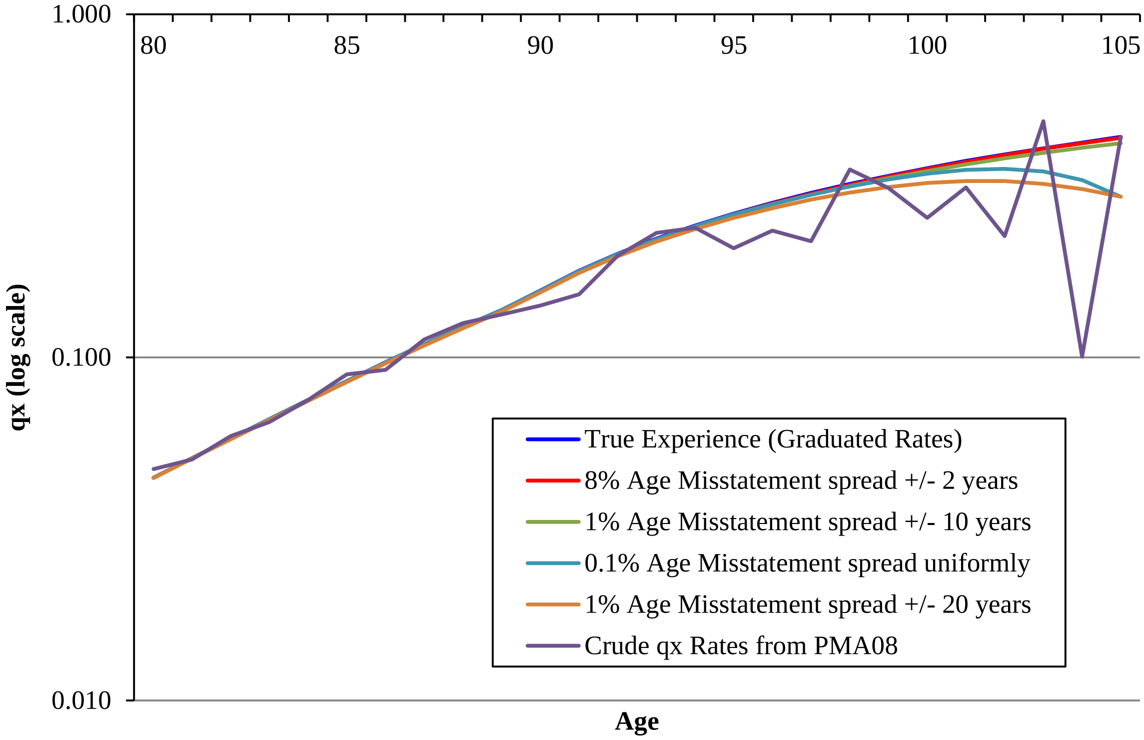

Table 6.8: Graduated Rates Life Expectancies versus levels of age misstatement

	Graduated Rates	8% Age Misstatement spread +/- 2 years	1% Age Misstatement spread +/- 10 years	0.1% Age Misstatement spread uniformly	1% Age Misstatement spread +/- 20 years
Period Life Expectancy					
65	20.22	20.23	20.23	20.23	20.25
80	8.95	8.96	8.97	8.96	9.01
90	4.10	4.12	4.13	4.13	4.20
Percentage difference versus Graduated Rates					
65	100.0%	100.0%	100.0%	100.0%	100.1%
80	100.0%	100.1%	100.2%	100.1%	100.7%
90	100.0%	100.4%	100.7%	100.6%	102.3%

It is clear from the examples above that the key driver of the impact of age misstatement is the number of years by which ages are misstated rather than the proportion of records affected by age misstatement. Relatively narrow age ranges hardly distort the curves, even if a reasonably large proportion of records are affected (see 8% +/- 2 year example). For this example, life expectancy at 65 is materially unchanged and reduced by 0.1% at age 80. With a large spread and a relatively small proportion affected (0.1% spread uniformly scenario), we see again relatively no change in life expectancy at 65 and a 0.1% reduction in life expectancy at age 80.

The most severe misstatement scenario we have considered (1% +/-20 years) has a 0.1% impact on life expectancy at 65, 0.7% at age 80, and 2.3% at age 90.

6.1.6. Interaction between delays in death reporting and age misstatement

We have attempted to illustrate the interaction between age misstatement and death reporting delays given both of these effects lead to a lowering of the mortality curve at higher ages. Figure 6.5 looks at a particular scenario of age misstatement and death reporting and compares the joint impact. We have found that the results of this analysis are not particularly sensitive to the order in which age misstatement and reporting delays are applied. The example illustrated in Figure 6.5 and Table 6.9 assumes age misstatement then reporting delays.

Figure 6.5: Interaction of age misstatement and late reporting of deaths scenarios: 1% Age Misstatement spread +/- 10 years; Constant 3.5% deaths unreported from 2007

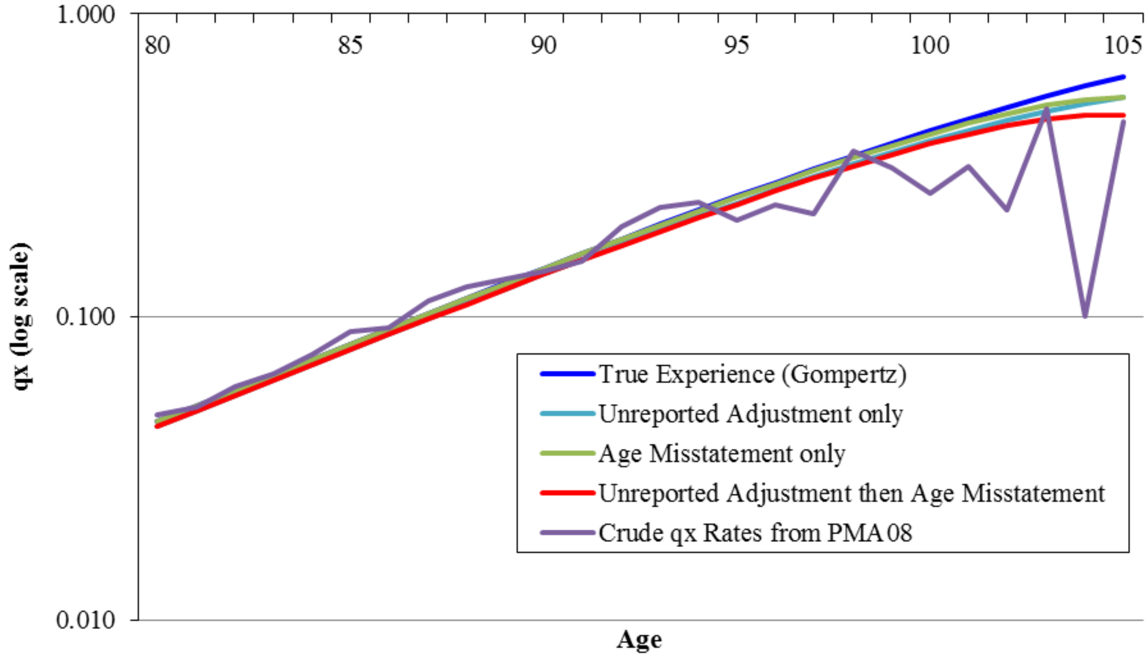

Table 6.9: Gompertz Life Expectancies versus levels of late reporting and age misstatement

	Gompertz Curve	Constant 3.5% unreported (from 2007)	1% Age Misstatement spread +/- 10 years	Late reporting of deaths then age adjustment
Period Life Expectancy				
65	20.33	20.63	20.34	20.64
80	9.17	9.39	9.18	9.41
90	4.31	4.49	4.34	4.51
Percentage difference versus Gompertz				
65	100.0%	101.4%	100.0%	101.5%
80	100.0%	102.5%	100.2%	102.7%
90	100.0%	104.0%	100.6%	104.6%

Figure 6.5 shows that the age misstatement effects applied after the late reporting of deaths adjustments seem to be relatively additive which is confirmed in the life expectancy calculations under the various scenarios.

6.1.7. Heterogeneity in high age data

Section 4.1 introduced the idea that mortality deceleration observed in portfolio and population data could be explained by heterogeneity in the underlying data being analysed. That is to say that the data is comprised of various homogeneous groups which individually exhibit mortality in line with a Gompertz law but when analysed concurrently lead to the derivation of a mortality curve which falls away from a Gompertz curve with increasing age.

The purpose of this section is to examine the extent to which this theory can be demonstrated using a model population and a real life example.

Heterogeneity in mortality level

To conduct this analysis we consider a starting population which exhibits mortality in line with the Gompertz curve derived in sections 6.1.3 and 6.1.4 using the crude data from the PMA08 graduations. We have then considered three populations, the central population exhibiting mortality in line with this curve and the other two populations in line with level adjustments around this curve.

In line with sections 6.1.3 and 6.1.4, we have assumed our three populations have the same number of members alive at age 65 and the annuitants alive at older ages are then derived from the assumed underlying mortality rates.

We have considered adjustments to mortality level of + / - 0.5 to log μ_x. This example was chosen as a relatively extreme difference and gives rise to differences in life expectancy at age 65 of around 7 years between the highest and lowest longevity groups. This compares to differences in life expectancy between rich and poor of 8 years according to Bennett et al (2015).

We have created combined mortality curves by summing the derived deaths and exposures for each age from each of the three populations for comparison with the underlying three curves. In Table 6.10 we set out the life expectancies for each group and the combined group as well as the differences in survival and q_x curves.

Table 6.10: Gompertz Life Expectancies versus adjustments to log µx and results on qx curves of combining all three populations

	Gompertz graduated q_x rates	Gompertz graduated q_x rates with -0.5 adjustment to log μ_x	Gompertz graduated q_x rates with 0.5 adjustment to log μ_x	Combined q_x curve
Period Life Expectancy				
65	20.34	23.94	16.93	20.41
80	9.17	11.81	6.89	9.54
90	4.31	6.01	3.00	4.98
Percentage difference				
65	100.0%	117.7%	83.2%	100.3%
80	100.0%	128.8%	75.2%	104.1%
90	100.0%	139.2%	69.5%	115.4%

Figure 6.6: Survival curves for 3 homogeneous populations with level differences relative to survival curve from combined experience

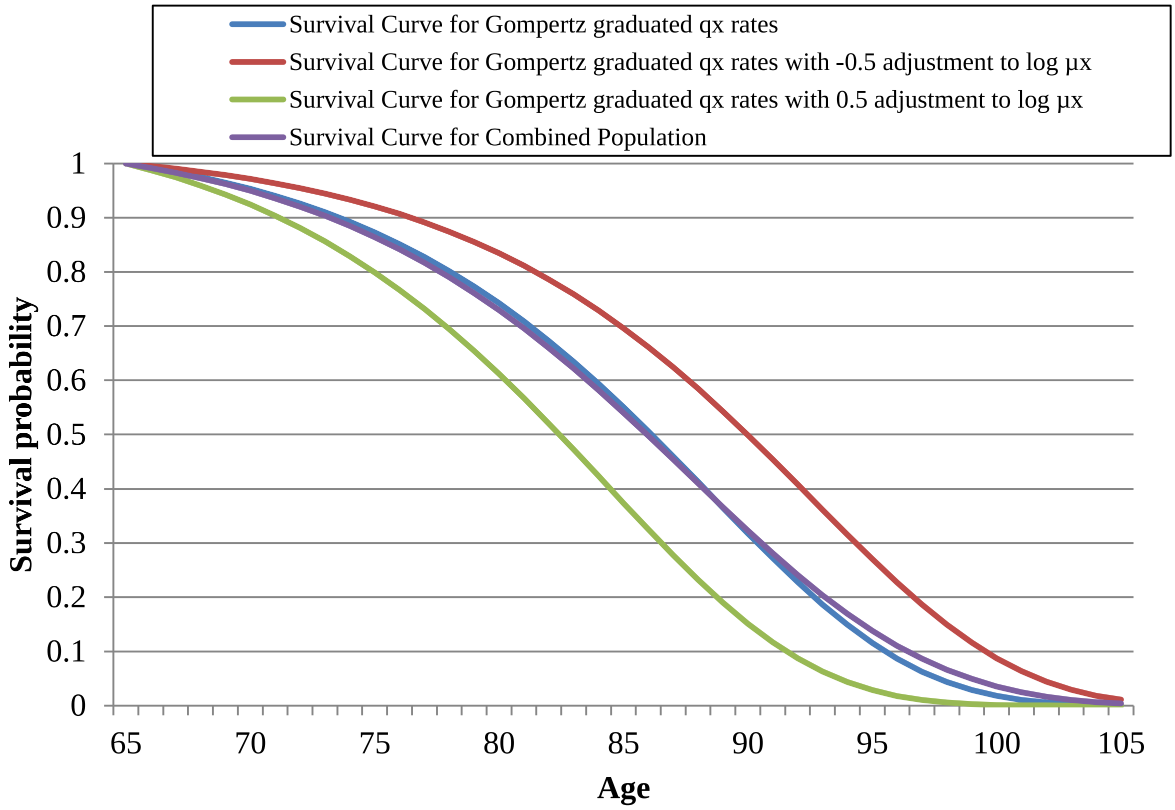

Figure 6.7: q_x curves for 3 homogeneous populations with level differences in μ_x relative to q_x curve from combined experience

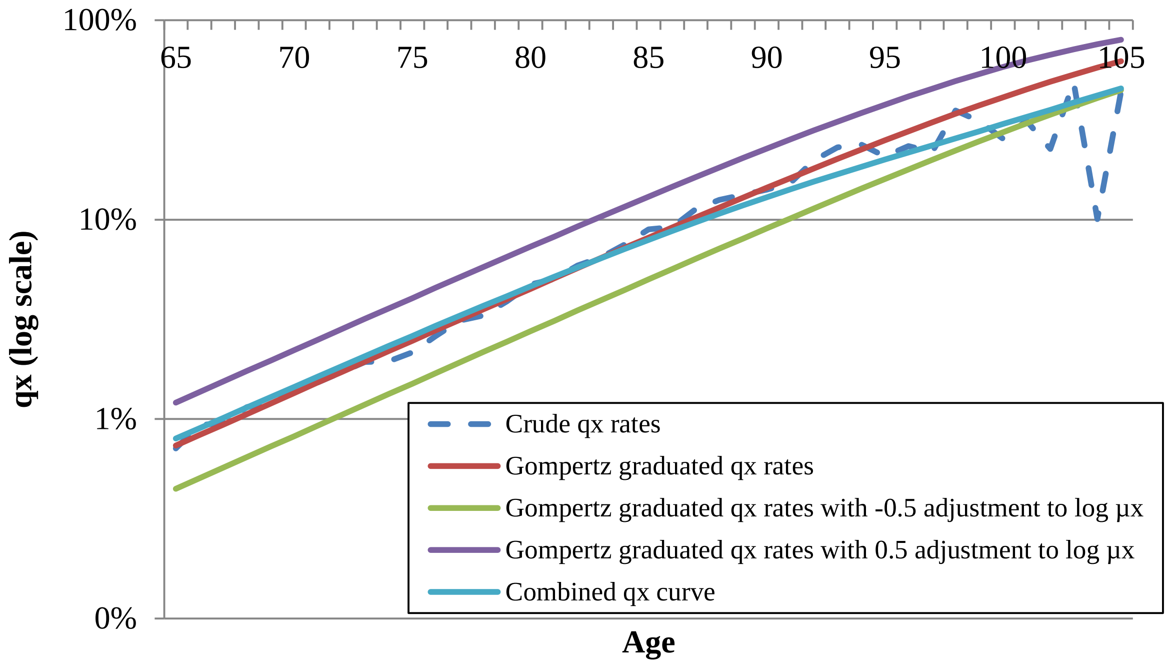

The survival curves in Figure 6.6 illustrate how the effect of combining the populations differs at the younger and older ages with the heavier mortality group dominating at younger ages (the combined curve sits below the central case / graduated rates) and the lighter group at older ages since most of the heavy group are deceased (the combined curve sits above the central case / graduated rates).

It is interesting to see in Figure 6.7 that the combined q_x curve fits the crude rates at older ages quite well, illustrating that heterogeneity in mortality level could explain some of the observed deceleration in the crude rates.

The impact on life expectancy at 65 of the combined curve versus the original curve is low, most likely due to the offset between differences at low and high ages. However, the impact at older ages is large, with a 4% difference at age 80 and 16% at age 90.

This section has demonstrated that heterogeneity in portfolio and population data is a plausible explanation for plateauing of the mortality curve even if the underlying process driving mortality of each homogeneous group follows a Gompertz law. Therefore, it would seem reasonable to allow for some level of deceleration in mortality curves since actuaries are not often able to isolate homogeneous groups when pricing / reserving.

6.2. Data and modelling issues for key datasets

6.2.1. UK population data

UK population data – exposed to risk

Gathering and publishing population data for the constituent countries of the UK is the responsibility of different bodies; data for England & Wales are published by the ONS, data for Scotland by National Records of Scotland (NRS) and data for Northern Ireland by the Northern Ireland Statistics and Research Agency (NISRA). Some data at a UK level is published by ONS but usually this can only be obtained by aggregating data published by the three separate bodies. There can be some differences in data or definitions between the bodies for both population estimates and for deaths; for example, arising from differences in the registration process. However, whilst this means that data from the different countries may not be available on entirely consistent bases, the differences are generally small for the purposes of calculating national mortality rates. (Note that for the analyses included in this paper, the Working Party has used population data for England & Wales only.)

As discussed earlier, exposed to risk for population data is usually calculated using population estimates. For the UK and constituent countries, the exposed to risk for a given age in a calendar year is usually taken as the age last birthday mid-year population estimate for that year. These, in turn, are based on census data which are rolled forward to future mid-years by adding in the births, subtracting the deaths and allowing for net migration. This gives rise to specific issues, some of which are discussed in Cairns, Blake, Dowd & Kessler, (2014).

Estimates of the population at mid-year by sex and single year of age last birthday up to age 89 and at ages 90 and over are published annually for each of the constituent countries of the UK. Additionally, estimates of the mid-year populations by single year of age last birthday for ages 90 and over are also published. These are produced using a different methodology from that used to roll forward the census data but the resulting age-specific estimates are constrained to add up to the rolled-forward estimates of the population aged 90 and over by sex. Published data are provided to the nearest ten but estimates to the nearest integer are available on request.

If the census population estimates are significantly different from the rolled-forward estimates then historical population estimates may be revised following a census.

It should be remembered that these figures are all estimates rather than counts. Whilst the registration systems in the UK mean that data on births and deaths are thought to be reliable, there are some cases where deaths may not be registered for several months or years (e.g. if an inquest is involved). There are no accurate measures of internal or international migration, so this introduces an element of uncertainty in the figures; in particular, even if there is relatively little international migration at high ages, migration between the constituent countries of the UK may have some effect if mortality rates for each country are calculated. Even if the movement data used to roll forward were all accurate, the census data itself is an estimate, not a count, and is subject to data misstatements, entry errors, imputations etc.

There are likely to be similar problems as for life office and pension scheme datasets with ‘phantom’ people being included in population estimates. If there are more people estimated at age x in one census than there really were in the population at the census date then, even if the data on subsequent deaths and migration were totally accurate, these people would never be removed from the database. (This is a slightly different issue from the phantoms in life office and pension scheme data which normally relate to people who did exist but were not subsequently removed from the database following death.) For the population data there can also be 'negative phantoms' if fewer people were estimated at the census date than were actually in the population. For such a case this could result in deaths at some old age in a given year that outnumber the number of people who could have died at that age in that year.

As well as the general issues discussed above, there are other specific issues which affect population data, including:

- The mid-year population estimate may be a poor approximation of the actual exposed-to-risk, particularly where there is a period of non-uniform births in a calendar year (for example, the 1918, 1919 and 1920 birth cohorts).
- Historical estimates may be revised following a new census; if this is done, the methodology used may still give rise to phantoms in the data in past years, as it is generally assumed that the previous census figures were correct.
- The rolling forward of the 2001 census figures to mid-2001 assumed an even spread of birthdays during each year at all ages. This is reasonable for years where birth rates were relatively stable over the year and from year to year. However, as mentioned above, there are some birth cohorts for whom this is not a good approximation. This was not an issue for rolling forward the 2011 census to mid-2011 where the actual incidence of birth was allowed for.
- As noted earlier, the method used by ONS for producing population estimates by single year of age at ages 90 and over is different from that used to produce population estimates at younger ages. This can produce a discontinuity in the estimates for ages 89 and 90. It should also be noted that the historical estimates at these ages are revised each year when the estimates for the following year are produced.

These issues are discussed in more detail in Cairns, Blake, Dowd and Kessler (2014).

UK population data – deaths

At a national level, deaths in England & Wales have to be registered within five days of occurring, except in certain circumstances, such as if a coroner or an inquest is required. Deaths resulting from external causes give rise to the longest delays (e.g. transport accidents). The situation in Scotland is slightly different as deaths must be registered within eight days and deaths can be registered even if a cause of death is not known at the time (this should lead to shorter reporting delays with Scottish data).

Hence, in theory, for population data there should be no problems arising from deaths not being notified, which might arise in life office or pension scheme datasets. However, there are other data issues similar to those for life office or pension scheme data as well as others which are specific to population data.

Deaths of those whose usual residence is outside England & Wales are included in the total figures for England & Wales but not in any subdivision of England & Wales. Deaths of Her Majesty's Forces occurring outside England & Wales are included whilst deaths to foreign

Armed Forces stationed in England & Wales are excluded. ONS mortality statistics exclude deaths of all residents of England & Wales that occur and are registered outside England & Wales.

Information collected on death registration in England & Wales is recorded on the Registration Online (RON) system by registrars. Most of the information is supplied by the informant (normally a close relative of the deceased). The date of birth is recorded rather than the age at death and, where required, age at death is derived from the data of birth and the date of death. Hence problems with age misstatement will only arise in cases where an incorrect date of birth (or possibly date of death) is provided or recorded.

The informant must visit the registrar's office in the district within which the death occurred. The registrar will ask the informant to check that the information is correct before the registration is submitted. Since April 1997 information can be given to a registrar elsewhere (registration by declaration) who will then send the information to a registrar in the district where the death took place. Information supplied by an informant is generally believed to be correct as knowingly supplying false information may render the informant liable to prosecution for perjury.

The registrar is instructed to make the following checks:

- the death is in their area;
- the death occurred in the last 12 months;
- the informant is qualified to give information;
- the correct medical certificate of cause of death has been used;
- the certificate relates to the correct person;
- the certificate has been filled in properly – it is signed, it has not been amended, it has the doctor's qualifications, the last date seen alive and whether the certifier saw the deceased after death; and
- the death does not need to be referred to the coroner.

The register page is signed by the registrar and the informant.

Quarterly checks are carried out on the entries that they are in sequence, there is a medical certificate/coroner's form to accompany each death entry, each entry has been signed by an informant (if required) and the registrar and a general check is carried out on any manual entries for illegibility, erasures etc.

When data is entered into RON, validation checks are carried out to help ensure the details entered are correct. Any issues identified are resolved with the individual registrar. The checks include:

- identification of missing entries;
- checks for duplicate records;
- checks for misplaced records;
- for paper records – that the date of death and date of registration are in the correct range and are complete prior to keying in; and

- checks on registrars whose returns are not received by the fourth working day after the end of each week.

Once on the ONS database, data are passed through a series of automatic validation processes which highlight any inconsistencies. Simple validations include examination of dates or employment status to ensure they are likely. More complicated checks include checks for consistency between dates of birth, death and registration or between age and marital status. Further checks are carried out before the annual mortality dataset is finalised; such as checking the underlying cause of death is compatible with the sex, frequency checks on a range of fields such as age, sex, underlying cause and area of residence.

Deaths data at the national level are published on a registration basis i.e. deaths are allocated to the year in which they are registered. This may be later than the year in which they occurred, for example for deaths which occur at the end of a calendar year but are not registered until early in the following year or, for England & Wales data, because a coroner or an inquest is involved. Deaths are also available on a year of occurrence basis. These are obtained by taking an extract from the deaths database at a given date; deaths which occurred before this date but which are not registered until a later date will not be included in this published occurrence data. Hence in general the dataset of registered deaths will contain all deaths that have been registered in a calendar year so historically are complete on a registration basis, but will contain deaths which occurred in earlier years. Deaths on an occurrence basis will contain deaths allocated to the correct year in which they occurred but will not be complete, at least for more recent years.

Note that for 1993 to 2005, deaths data in England & Wales were published on an occurrence basis rather than a registration basis so there are some issues in using the time series of published deaths data for England & Wales; data on deaths on a registration basis are available from ONS for these years.

Table 6.11 shows the numbers of deaths registered in England & Wales for 2001 to 2014, the percentage of those registered which occurred in the same year and the number of deaths which occurred in the year (note that the latter is not a complete count of occurrences in each year). As can be seen there is relatively little difference in the numbers of deaths registered in a year and those that occurred in a year (for most years the difference is less than 0.5% – this is also the case just considering deaths at ages 95 and over). The percentage of registered deaths which occurred in the same year has been gradually falling over time although it is not clear why this should be happening.

Table 6.11: Deaths by year in England & Wales on a registration and occurrence basis

Year	Deaths registered in year	Percentage of registered deaths which occurred in same year	Deaths occurring in year (not complete)
2014	501,424	95.3%	N/a
2013	506,790	95.2%	N/a
2012	499,331	95.9%	499,240
2011	484,367	95.7%	481,156
2010	493,242	96.0%	491,449
2009	491,348	95.9%	n/a
2008	509,090	96.0%	n/a
2007	504,052	96.2%	n/a
2006	502,600	96.5%	n/a
2005	512,993	97.0%	512,692
2004	514,250	97.1%	512,541
2003	539,151	97.3%	538,254
2002	535,356	97.3%	533,527
2001	532,498	97.1%	530,373

Source: ONS

Of the 501,424 deaths registered in 2014, 477,752 occurred in 2014, 21,474 occurred in 2013, 1,655 occurred in 2012 and 543 occurred prior to 2012.

An analysis by ONS of deaths registered in 2011 showed that around 78% were registered within five days and 94% within one month of the date of death. Delays in registration varied by cause of death, the median delay was around two days for some forms of cancer but 139 days for deaths from external causes. No information was provided broken down by age at death.

Assuming a correct estimation of the relevant exposed to risk, the use of registration deaths data will give mortality rates which may be higher or lower than the actual mortality rate as not all the deaths registered occurred in that year; the use of occurrence deaths data will underestimate the mortality rate as not all the deaths which actually occurred in the year would be included in the published data.

Deaths which occurred in England & Wales should all be registered by law; this includes deaths of non-England & Wales residents (for years for which a split has been published, deaths of overseas residents accounted for between 0.2% and 0.3% of all deaths in England & Wales on an occurrence basis). However, as mentioned earlier, deaths of England & Wales residents who died outside England & Wales are not included in the data. Hence, the deaths data do not always correspond to the exposed to risk as some deaths of England & Wales residents are not included, whereas deaths of non-residents who died in England & Wales are included.

Some deaths may be allocated an incorrect age, particularly where the informant is not a relative of the deceased (for example if the deceased had no family) and the correct date of birth is not known.

Another issue that can cause problems is that deaths data are published by deaths in a calendar year by age last birthday at death, whereas the population estimates are published by age last birthday at the mid-year. Hence, the age definitions for the exposed to risk are not the same as those for the deaths.

6.2.2. The Human Mortality Database

The Human Mortality Database (HMD) contains mortality rates and life tables for various countries and the original raw data (births, deaths, census counts or population estimates) from which they were derived.

Various adjustments are made to the raw data to derive data in the same format for each country as far as possible by single year of age to 109 and 110 and over. Hence, it should be remembered when using data from the HMD that these may have been obtained by adjusting raw data, particularly at the oldest ages.

For computing mortality rates and life tables, further adjustments may be made as follows:

- For the deaths data collected, adjustments are made to distribute any deaths of unknown age proportionately across the age range and to split aggregated deaths into finer categories. Deaths provided on a calendar year basis are also split into deaths within a year by year of birth using a prescribed methodology.
- Population estimates below age 80, as at 1st January, are either taken directly from official estimates or derived using intercensal survival. Above age 80, population estimates are derived by the method of extinct generations for all cohorts that are extinct and by the survivor ratio method for non-extinct cohorts who are older than age 90 at the end of the observation period. Non-extinct cohorts aged 80 to 90 at the end of the observation period are obtained either from official estimates or by a method of intercensal survival.
- Exposed to risk during an age-time interval is based on 1 January population estimates with a correction that reflects the timing of deaths during the interval.

Further details on the methodologies adopted by the HMD to adjust and clean data are available at http://www.mortality.org/

6.2.3. Data and modelling issues for CMI datasets

This section summarises how the data and modelling issues set out in section 6.1 affect the CMI SAPS and Annuities datasets, as described in CMI working papers.

Estimation of exposure

The SAPS investigation uses member by member data (including dates of birth, death and retirement as supplied by the actuarial consultancy that submitted the data). Exposure periods and ages (for both exposure and deaths) can therefore be calculated using exact day counts rather than census methods.

The data underlying the Annuities Committee's most recent analyses ("08" Series tables) is primarily at an individual policy / benefit level, enabling calculation of exposure periods on a day count basis.

Derivation of the force of mortality

The SAPS graduations use central exposed to risk (calculated using exact day counts as noted above), and assume that deaths follow a Poisson distribution. First, the force of mortality is fitted to the data. Then, initial rates of mortality are derived from the fitted force of mortality using the formula:

$$q_x = 1 - e^{-\int_0^1 \mu_{x+t}\,dt}$$

and approximating the integral using Boole's rule, i.e.

$$\int_0^1 \mu_{x+t}\,dt \approx \frac{7\mu_x + 32\mu_{x+¼} + 12\mu_{x+½} + 32\mu_{x+¾} + 7\mu_{x+1}}{90}$$

The "08" Series use a consistent approach, noting that those graduations fit central rates of mortality, and it is then assumed that m_x equates to $\mu_{x+½}$.

Delays in death reporting

The SAPS investigation deals with delays in death reporting by excluding data relating to the period 30 days prior to the extract date. CMI Working Paper 65 includes a description of an investigation into the impact of this assumption, and concludes that there was no evidence to suggest that extending the 30 day exclusion period would give significantly different results.

The data underlying the "08" Series tables is assumed to have negligible under-reporting of deaths as it relates to deaths incurred in 2007-2010 and reported by 31 December 2011.

Duplication

The data underlying the "08" Series tables has very limited de-duplication applied. Offices supplying data were asked not to remove duplicate policies; however, it was not possible to combine duplicate deaths and exposures for lives with multiple policies across offices or even – in many cases – within an office.

7. Mortality of closed cohorts

7.1. Pattern of high age mortality by age

We have assessed the accuracy of the mid-year population estimates for England & Wales, as reported by the ONS and the HMD. This enables us to further assess the impact of data issues and modelling assumptions on population mortality.

We have considered registered deaths for cohorts of older lives for the England & Wales population where we believe the cohorts are essentially extinct (i.e. all lives in the cohort have died); we refer to these as "closed cohorts". To that extent, we have looked at the deaths registered for all cohorts in the England & Wales population that have reached at least their 110^{th} birthday. Cohorts that have not yet reached their 110^{th} birthday are deemed to be incomplete in terms of mortality records and have not been considered further, to avoid making assumptions about the mortality of incomplete cohorts.

Where a cohort is known to be completely (or nearly) extinct, we can estimate populations directly from recorded deaths. This provides a different estimate of population than using ONS population estimates which are linked to census information. The latter approach quickly becomes out of date, as census surveys are only performed every 10 years, leading to retrospective adjustments to population estimates.

Past population estimates for different cohorts are estimated by recursively adding back registered deaths at each age. For a given cohort aged x in calendar year y, we have used the following formulae:

- Population for age x in calendar year y = $P_{x,y}$
- Deaths for age x in calendar year y = $D_{x,y}$
- $P_{\max,y} = D_{\max,y}$, where max reflects the highest age banding used – for the ONS this is age 125+, for the HMD this is age 110+
- $P_{x,y} = P_{x+1,y+1} + D_{x,y}$.

The extinct generation approach assumes:

- The recording of deaths is accurate and complete. We note however that the data is subject to the quality issues set out earlier in this paper.
- There is no international migration. We believe net migration for England & Wales is low at old ages, and we expect migration to be negligible at very old ages.

Key data issues include territorial changes over time, significant war mortality and estimation of military population, and changes in recording and validating deaths over time. These have not been addressed in our investigation, but may be considered for further analysis.

The extinct generation method returns the population at the start of each year and hence naturally provides q_x rates. We have also determined crude m_x from ONS data, and converted those to q_x rates using the approximate formula: $q_x \approx m_x / (1+m_x/2)$. This is necessary for comparison purposes, as the ONS population estimates are at mid-year. This approach enables us to compare the mortality estimates on a like-for-like basis.

We present below mortality rates as implied by closed cohorts under the extinct generation method approach relative to mortality as implied by the death and mid-year population estimates as provided by the ONS and HMD. We have presented mortality for males and females separately, with heat maps showing the relative mortality differences for each age between 60 and 95 for the period 1961-2003. We have focused on ages up to 95 to remove the increased volatility observed from conducting an extinct generation approach to very small numbers.

The ONS and HMD record (almost exactly) the same numbers of deaths by gender and by individual age. However, the ONS and HMD apply different approaches to determining the mid-year population estimates at very high ages (in addition to slightly different approaches overall):

- The ONS currently publishes population estimates by single year of age up to age 90, with a final category for ages 90 and above. They then apply a Kannisto-Thatcher approach to effectively consider the ratio of the number of survivors belonging to a cohort still alive against that of a cohort where the lives have died in the last few years. By making an assumption about the maximum age that everyone in a cohort lives to, the ONS then apply an algorithm to apply survival ratios that will return estimates of the number of lives in each cohort alive at earlier ages. These estimates are then constrained to the official mid-year population estimate of the total England & Wales population aged 90 and over. The over-90s estimates are revised each year.
- The HMD takes a different approach to modelling population exposures above age 80, as follows. For all cohorts that are extinct, population estimates are determined by the method of extinct generations. For non-extinct cohorts who are older than age 90 at the end of the observation period, the survivor ratio method is used. To reflect that different cohorts may observe different mortality trends, there is a further scalar adjustment to ensure that the official population at age 90 is reproduced by the survivor method. For non-extinct cohorts aged 80 to 90 at the end of the observation period, population estimates are obtained by applying the method of inter-censal (or post-censal) survival. Acknowledging that recent mortality rates have been reducing and that survival ratios change over time, an adjustment is applied so that the official estimate of the population size for age 90 and above is maintained. This last feature means that for non-extinct cohorts the level of the population is the same as the ONS but that the shaping by age will be different.

Note that, in its analysis, the Working Party sought to replicate the application of the HMD methodology to ONS data but we have not yet been able to fully reconcile this; as a result the conclusions set out in this section should be regarded as provisional.

Both the ONS and HMD approaches apply a form of survival ratios. The HMD approach applies extinct cohort mortality direct from the death information where all lives belonging to a cohort are known to have died.

Figure 7.1: Comparison of extinct generation mortality against ONS estimates

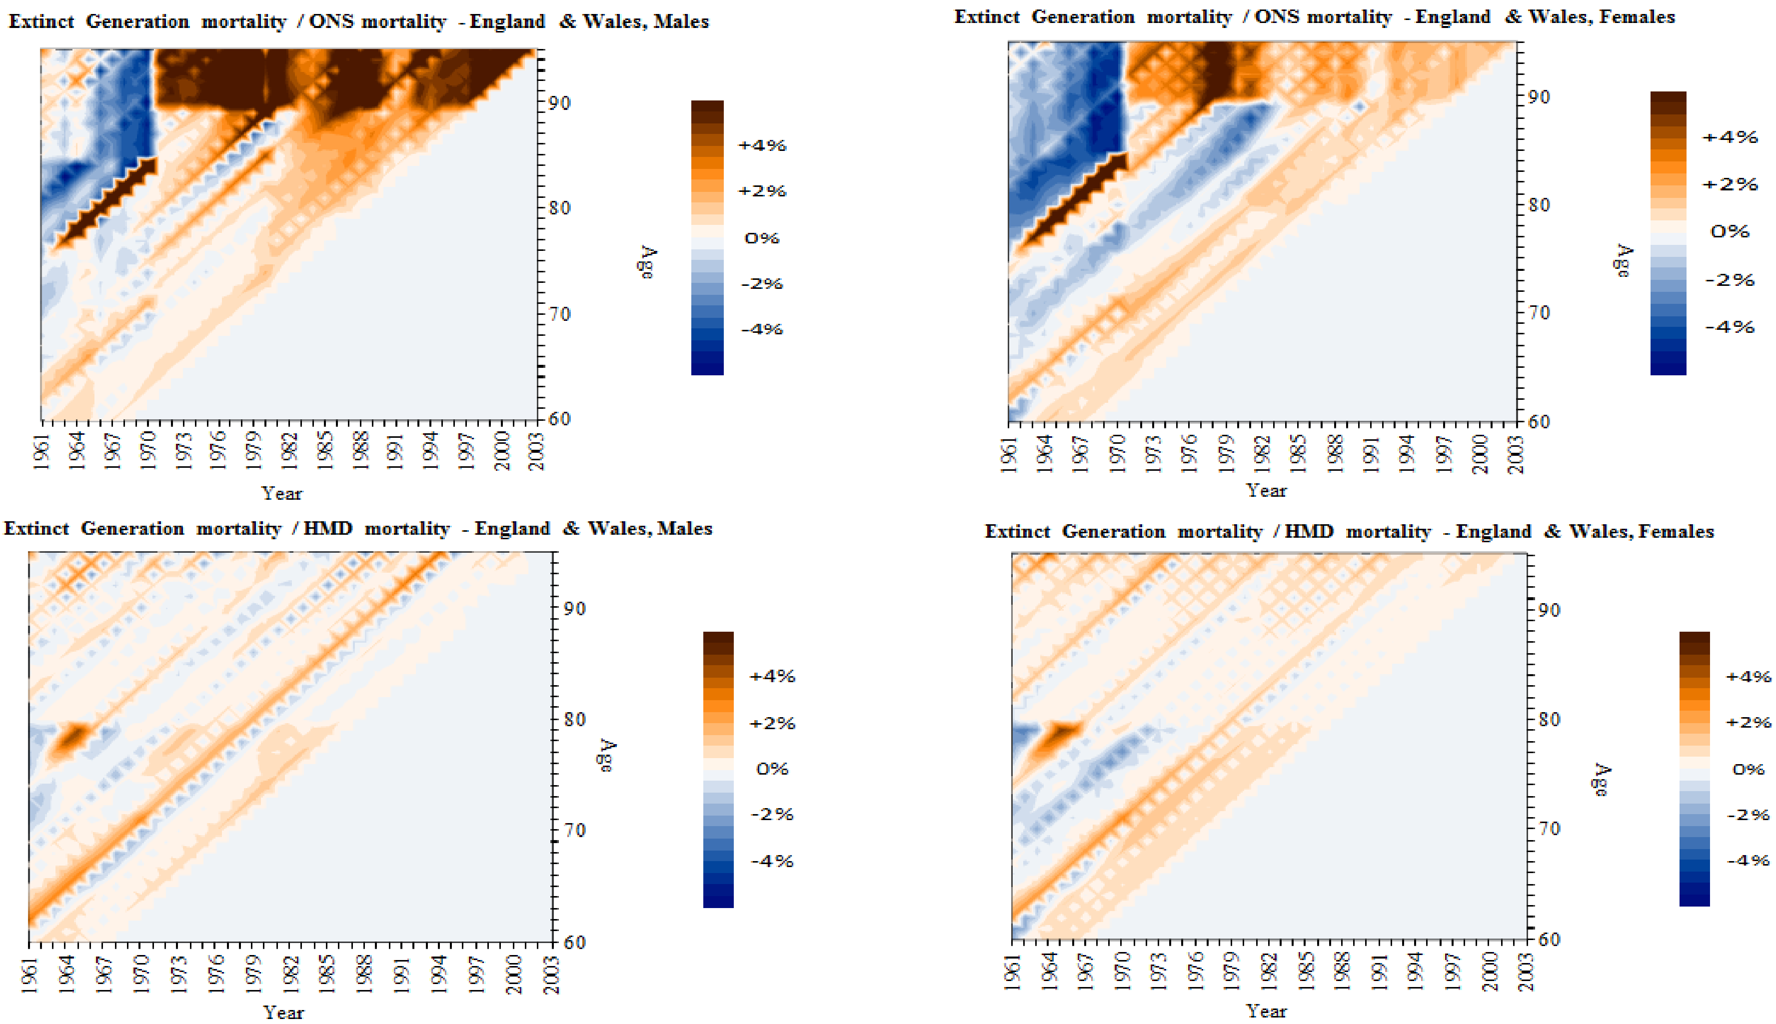

The four graphs in Figure 7.1 compare the historical mortality rates for England & Wales determined using the extinct generation approach against the ONS and HMD approaches for males and females. Warmer colours (brown/orange) represent higher extinct generation mortality rates, whereas colder (blue) represent higher ONS (or HMD) mortality. The pale blue triangle in the bottom right of each graph represents cohorts that have not been considered as their cohorts are not yet extinct.

We can observe a number of key features from these graphs:

- There are strong cohort features, as demonstrated by the diagonal cohort lines.
- ONS mortality is higher than extinct generation mortality for ages 85+ prior to 1969, and lower after this point. The CMI Mortality Projections Committee reported on this age range's data, including data errors that have not been retrospectively corrected, and excluded calendar year data prior to 1974 from the CMI_2014 version of the CMI Mortality Projections Model for this reason. This analysis confirms the data anomaly, and supports the decision to exclude this data.
- Broadly speaking, for recent years as age increases the extinct generation returns higher mortality and the gap between extinct generation and ONS increases with age.
- The difference between the approaches is more pronounced for males above age 90 than females. In recent years, mortality for ages 90 and above is typically greater than 5% above ONS, whereas for females it is typically around 1-2%.
- The difference between extinct generation and ONS approaches tend to increase during each decade (i.e. the colours show larger differences late in the decade than they do early in each decade). This is consistent with census inaccuracies underlying the ONS data increasing (in relative terms) as census data is projected over the course of each decade.
- We can observe that the wider differences observed in the ONS heat maps are not present in the HMD heat map analysis, confirming that the HMD approach adopted is a closer fit. We would expect the HMD mortality rates to be close to our own analysis given the HMD's approach to using a similar approach for extinct cohorts. There do remain some small cohort differences, as shown by the diagonal lines.

In their paper 'Phantoms Never Die' (2014), Cairns, Blake, Dowd and Kessler have assessed the level and shape of high age population mortality within the ONS data. They explored the use of concavity functions to describe how consecutive high ages deviate from linearity. If the log of death rates is linear then the concavity function should stay close to zero with no systematic bias exhibited either above or below. They concluded two potential features in population estimates. The first is that the mid-year population estimates for cohorts born immediately after the end of the First World War do not appear to be reasonable estimates for exposures given the non-uniform birth profile that occurred for these lives. The second was that the Kannisto-Thatcher approach employed to estimate populations for individual ages above 90 does not fit smoothly to individual age population estimates.

Both the Cairns et al and our own analyses indicate that there would be merit in considering the accuracy of the ONS population data further, with a particular focus on the mid-year population estimates as a proxy for exposure. This may result in recommended revisions to population estimates. Once this work has completed, we intend to then consider the historical mortality trends at high ages, and the implications for the CMI Mortality Projections Committee in its modelling of mortality improvements.

7.2. Time trends in high age mortality for closed cohorts

7.2.1. Background

The purpose of this analysis is to look at available data and investigate whether evidence exists of trends in mortality at very high ages. The analysis is separated out into nonagenarians (lives aged 90-99) and centenarians.

Method

In a similar vein to the previous section, we have used deaths registration data from the ONS for closed cohorts to construct l_x tables. In contrast to the previous section, we have assumed that cohorts that have not yet reached their 116th birthday are not closed and so have not been considered in this analysis.

We have constructed l_x tables starting from age 100 and also starting from age 90 by working backwards from the deaths data, in a similar way to the approach followed in section 7.1.

Summary of data

We have used the ONS deaths data by single year of age at death for England & Wales covering the period 1963 to 2013. We have 146,634 male and 526,527 female deaths aged 90 and above spanning the 15 birth-cohorts to 1898. We have 5,333 male and 37,897 female deaths aged 100 and above spanning the 25 birth-cohorts to 1898.

Both centenarian and nonagenarian populations have been steadily increasing, particularly females, as can be seen in Figures 7.2 and 7.3 which set out the number of lives reaching age 100 and age 90, respectively, from the oldest to the youngest birth cohorts.

Figure 7.2: Number of lives who reached age 100 in birth cohorts from 1874 to 1898

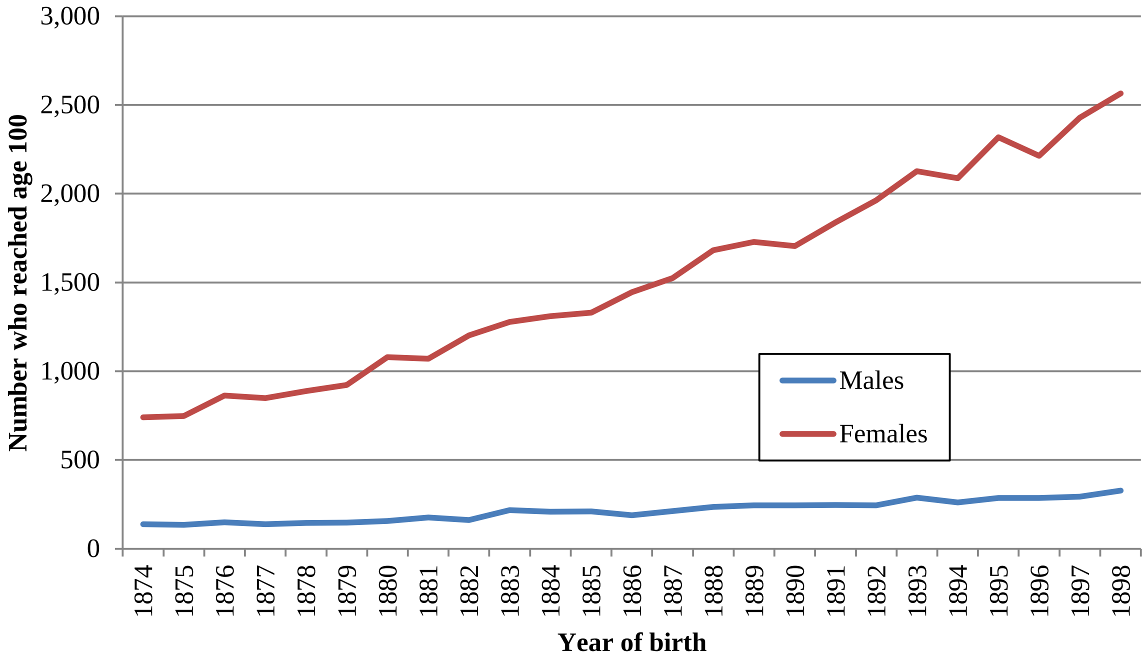

Figure 7.3: Number of lives who reached age 90 in birth cohorts from 1884 to 1898

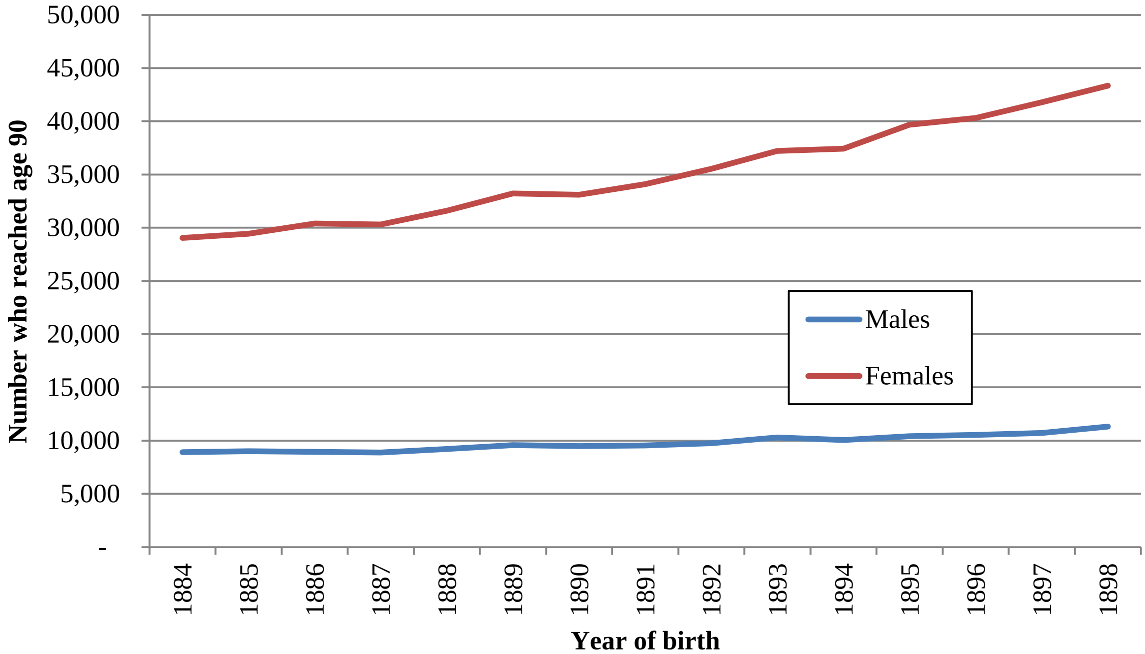

7.2.2. Median age at death

To investigate whether mortality rates at the highest ages have changed over time, we have looked at the median age at death through time for those who have reached age 100 and age 90 respectively within the closed cohorts being considered. By comparing the progression of median age at death through time for both groups we can observe whether mortality rates are changing and compare any differences between the 90+ and the 100+ groups.

We assume a uniform distribution of deaths between two integer ages. To calculate the median age, we have interpolated between the relevant ages such that the median life sits between the respective l_x and l_{x+1}. For example if 1000 people died over age 100 (l_{100} =1000), the median life would be life 500. If l_{101} =700 and l_{102} = 400, the median age at death would be 101.67.

We have considered the median age of death for centenarians and for those reaching age 90 and set out the results in Figures 7.4 and 7.5, noting that we consider a lower number of birth cohorts for those reaching age 90.

Figure 7.4: Median age at death for those who reached age 100 in birth cohorts from 1874 to 1898

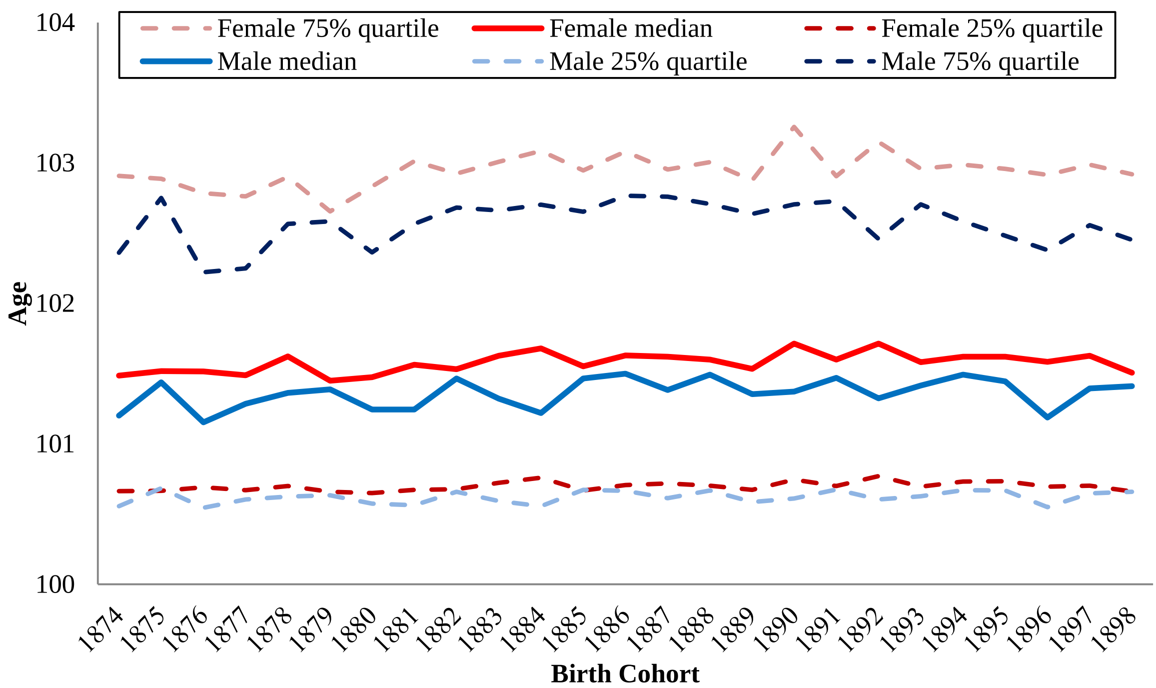

If we compare the average of the figures for 1874-1885 to the figures for 1886-1898 we see that both male and female median age at death have increased by 1 month. From this it does not appear that centenarians have seen a significant improvement in survival probabilities after age 100.

We now turn our attention to nonagenarians. In Figure 7.5, we see a more stable trend due to the larger number of lives in the study. The median age at death for males increases by around 3 months and for females by around 6 months, over the 15 birth cohorts analysed.

This could mean that mortality improvements in ages under 100 have been higher than those over age 100.

Figure 7.5: Median age at death for those who reached age 90 in birth cohorts from 1884 to 1898

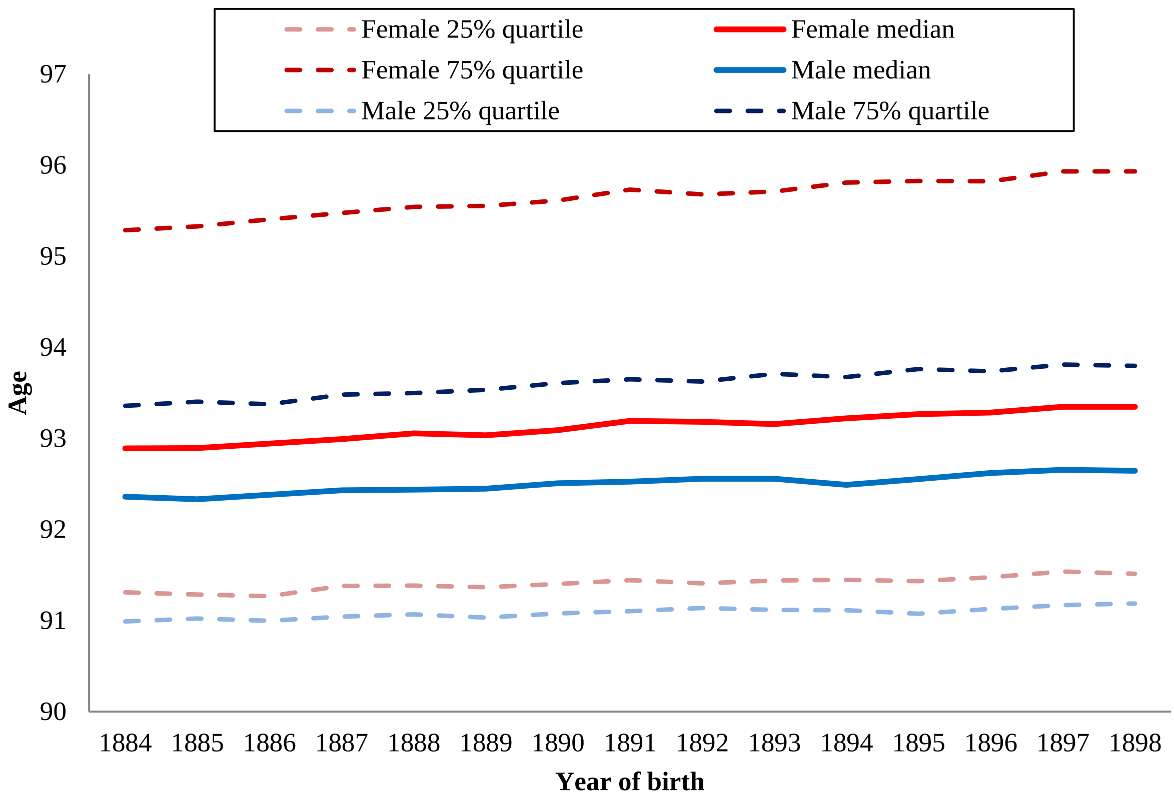

7.2.3. Probability of Survival

As an alternative way to look at how mortality rates at older ages have progressed over time we have analysed the probability of survival from ages 100 to 105 and ages 90 to 105 for the closed cohorts under consideration. In order to quantify confidence intervals around these probability of survivals we have assumed that the number of survivors X for a given birth cohort to age 105 follows a Binomial distribution with number of trials l_{100} and probability of survival $p \approx \frac{l_{105}}{l_{100}}$.

The standard error of the survival probability can therefore be estimated as:

$$\sqrt{\frac{p(1-p)}{l_{100}}}$$

The 5-year survival probabilities after age 100 (determined directly from the extinct generation death rates) are shown in Figure 7.6 together with the 95% confidence intervals. The interval for males is wider than for females because there are fewer males at the starting age 100.

Figure 7.6: Probability of survival to age 105 from age 100

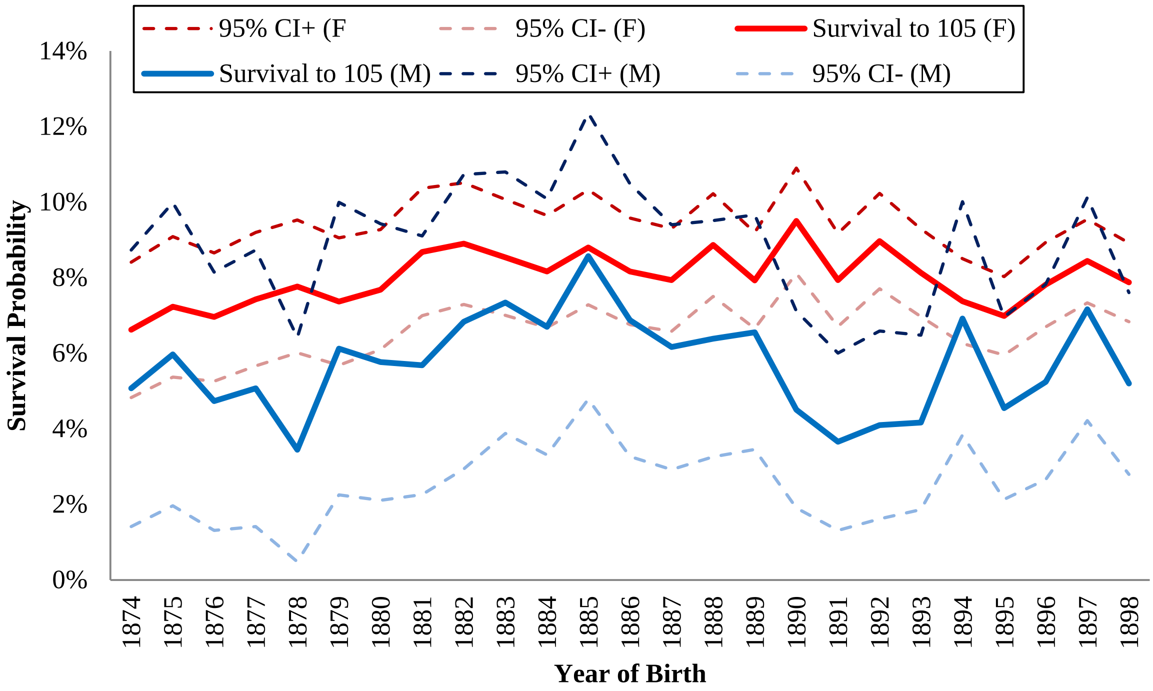

In line with the analysis of median age at death, it is difficult to conclude that the survival probability has changed much between birth cohorts. Figure 7.7 sets out the probability of surviving from age 90 to 105.

Figure 7.7: Probability of survival to age 105 from age 90

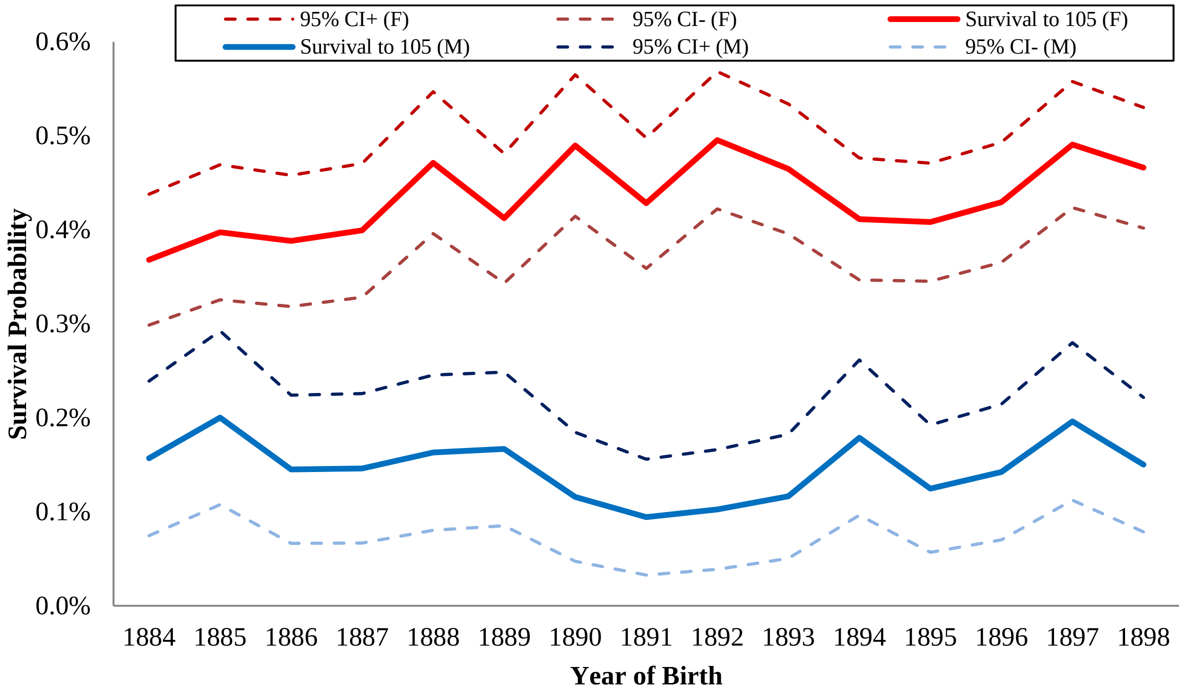

For each of the 15 cohorts analysed, the chance of a 90 year-old reaching age 105 does not appear to have improved for males whilst there seems to have been a marginal increase for females, with the increase in survival probability being of the order of 0.1%.

7.2.4. **Conclusions**

These results show no clear evidence for a material change in mortality patterns for England & Wales centenarians born between 1874 and 1898.

For those who have reached age 90, we see modest improvements in the median age at death and the survival probability from age 90 to 105 for females. For males we see a smaller increase in the median age at death and broadly no change in survival probabilities.

It should be appreciated that a simple model has been used here. For example, the survival analysis combines mortality across a range of ages, and as such may not reflect the changing shape of mortality by age. The Working Party is considering exploring time trends building on its work in section 7.1 as part of further activity in this area.

8. Do mortality rates converge with age?

Datasets often show heterogeneity at younger ages expressed as differences in absolute levels of mortality. We now consider the issues that readers may wish to take into account when deciding how far to extend this heterogeneity into the oldest ages.

Distinction needs to be made when drawing conclusions about convergence between (heterogeneous) "populations" and "individuals" as we can have divergence in relative mortality between individuals yet still see convergence with age for the population as a whole.

For example, in section 6.1.7 we considered a hypothetical example where the combined experience of three homogeneous populations with mortality level differences at all ages *converged* to the lightest mortality group curve. This demonstrated that observing convergence could be driven in part by the survival of lives who are the least frail, despite starting with heterogeneous populations with no old age convergence.

We now consider the theories for and against this.

Arguments for convergence

We cite possible reasons for convergence as provided by Hoffmann (2008):

1. Biological ageing acts as a leveller of social differences because biological processes assume dominance over social determinants and eventually everyone must die regardless of social circumstance.
2. The welfare state reduces old age socio-economic inequality through benefits and policy.
3. The effects of experiences from earlier life (i.e. working conditions) fade out.
4. The observed mortality differences get smaller in old age but only on the aggregate level because the surviving population is more homogeneous than in the younger ages where there is more unobserved heterogeneity - selective mortality.

Arguments against convergence

Hoffmann also gives arguments against convergence. These seem plausible too:

1. The effect of past unhealthy experiences is postponed until older ages.
2. There is a mutual accumulation of disadvantages in health and social status resulting in a "health stock".
3. Vulnerability increases in old age and makes differential exposures more harmful.

Results from different datasets

In Appendix C we investigate whether we observe convergence in different datasets by (1) socio-economic circumstance or levels of deprivation, (2) pension size, (3) health status at retirement, (4) lifestyle choices and (5) education.

Summarising these results we see:

1. Social deprivation and socioeconomic circumstance – little evidence for convergence by social deprivation but some mortality crossover. Data by socioeconomic circumstance shows convergence.
2. Pension size – clear evidence of convergence by pension size with some crossover.
3. Health status at retirement – clear evidence of convergence by health status with some crossover.
4. Lifestyle choices – evidence of convergence between smokers and non-smokers.
5. Education – conflicting evidence for convergence.

Conclusion

The majority of datasets we have presented show convergence. If we believe the population becomes more homogeneous with increasing age then we should consider the empirical results observed at these oldest ages.

On the other hand some may believe there is still some underlying heterogeneity even at the oldest ages. This point of view would not necessarily contest any observed convergence as that would be explained by the observed rates quickly becoming dominated by the experience of the lives with the lighter mortality. If they used these converged rates they would expect to overstate life expectancy for the heavier mortality lives that happened to be alive at the oldest ages as section 6.1.7 demonstrates.

This dilemma suggests we should avoid choosing a run in age that is too low otherwise heterogeneity can distort the results. By run in age we mean the age at which we start to converge to the level of the population mortality or a formula based extrapolation. Longevity actuaries might consider graduating to their own data as high up the age curve as possible in preference to using a hypothetical curve providing they have sufficiently credible data. Section 5.4 discusses options here.

9. Implications for modelling and graduating high age mortality

We hope that this paper serves as a starting point with which to consider the implications for modelling and graduating high age mortality. The paper has focused on collating recent and past research for this topic to provide an initial point from which we intend to expand our modelling.

The key findings from our analyses to date are:

- There are various issues with the data quality for all the data sources which might be used to model mortality at the oldest ages. It is not immediately clear how much the differences in results might be due to data issues. As such, it is difficult at this point to make any recommendations on a preferred approach or approaches to modelling mortality rates at the oldest ages. However, the differences in the resulting mortality rates derived using different models are generally not material except at the very oldest ages. Hence, it is unlikely that any one approach would produce results which are materially different from another except at the oldest ages. Given this and the data issues mentioned above, the Working Party does not feel that the choices made to extrapolate mortality rates at the oldest ages in recent graduated CMI tables were unreasonable.
- Analyses of historical mortality from extinct cohorts imply that mortality for the England & Wales population above age 90 as published by the ONS has been underestimated by around 5% for males and about 1-2% for females in the period considered. We understand the underestimation of mortality to be driven by overestimation of population exposures at very high ages. The Working Party intend to consider this feature further in the next phase, including the potential impact on both current (base) levels of mortality and on the use of revised estimates on future mortality projections. Please see section 7.1 for further details.
- There is wide variation in the level and shape of mortality assumed at high ages under different tables published by the CMI, the ONS and North American actuarial associations. If the ranges of methodologies typically adopted are applied to the data underlying the S2PML tables then we observe a variation of -4.1% to +0.4% in cohort life expectancy (or -3.7% to +0.3% in annuity value with a 3% discount rate) for a male aged 90. Impacts are smaller at age 65, being -0.4% to +0.1% on cohort life expectancy and -0.2% to +0.0% on annuity value. Please see section 5.5 for further details.
- The debate on whether the shape of mortality at high ages is exponential or exhibits signs of mortality deceleration is inconclusive. The S2PML tables effectively assume mortality deceleration occurs. The impact on cohort life expectancy for a male aged 90 of instead adopting a Gompertz mortality shape at high ages is -2.5% to -0.5% (or a -2.0% to -0.4% impact on annuity value). Again, impacts are smaller at age 65, being -0.4% to -0.1% on cohort life expectancy and -0.2% to -0.0% on annuity value. Please see section 4.2 for further details.
- We have modelled scenarios considering the impact of late reporting of deaths at high ages. The potential impact on annuity values and life expectancies for a male aged 65 if late reporting is understated is typically in the region of +1% to +2% (and can be as high as +5% to +10% at age 90 in some of the scenarios considered). The delay in

reporting deaths has a more material impact at very high ages where the higher rate of mortality will have a bigger impact on restating exposures once all deaths are known. Please see section 6.1.4 for further details.

- We have modelled scenarios considering the impact of age misstatement. This suggests the potential impact on annuity values and life expectancies for a male aged 65 from age misstatement is in the region of 0.0% to +0.1% (and up to +0.5% to +1.0% for a male aged 90). We have considered the impact of date of birth corrections from data cleansing of large buy-in clients of a bulk annuity provider in constructing these scenarios. The impact varies by the range of age misstatements. Please see section 6.1.5 for further details.

We hope that the areas explored in this paper provide some useful reference for readers considering high age mortality. Each reader will need to consider how each of the areas explored in this paper might impact their own portfolio or population of lives given their quantity and quality of data, age profile, and product terms, as well as the impact of applying expert judgment in this area.

Areas of future research are set out in section 10 of this paper.

10. Future research

The High Age Mortality Working Party intends to explore further modelling impacts. These include:

- Population mortality at high ages: we intend to consider further the appropriateness of estimated mortality for the England & Wales population. In their 2014 paper, ‘Phantoms Never Die: Living with unreliable mortality data’, Cairns et al indicate issues associated with the ONS re-stating their point estimate for population exposures above 85 across all high ages using extinct generation methodologies. There are also concerns around the mid-year population estimates for particular cohorts, most notably for lives born just after the end of the First World War. We intend to explore these areas further and consider the appropriateness of the ONS approach and whether there are more appropriate alternative approaches. We will also explore whether there are any other datasets, such as longitudinal studies, which may provide additional useful information.
- Mortality trend: this paper focuses on issues that may affect historical and recent levels of mortality. We intend to explore how the analysis described above might help inform a time analysis of mortality trends, in particular for assisting the CMI Mortality Projections Committee with their modelling at high ages.
- International comparators: we also intend to extend our analyses to consider other territories and consider where similar features, or otherwise, are observed relative to the England & Wales population.
- Mortality convergence: we intend to consider the case for mortality convergence between groups of lives at very high ages.
- Mortality table construction: we recognise that this topic is of interest across the CMI committees. We intend to consider proposing approaches to graduating assurance, annuitant and pensioner mortality rate tables for the oldest ages.
- Seasonal mortality: we intend to consider the seasonal shape of mortality at the very high ages and how this differs relative to other ages.
- High Age mortality datasets: we intend to consider working further with CMI committees on the data issues and modelling features raised within this report, namely late-reporting and age of death reporting.

Our intention is to produce a follow-up paper considering these themes in 2016.

We welcome feedback both on this paper and our intended plans. We can be contacted at HighAgeMortality@cmilimited.co.uk.

References

A complete list of references can be found below. A Reading List, providing details of additional papers that were reviewed by the Working Party in their research, is published alongside this paper, should readers wish to explore this area further.

Bale and Lu, “Mortality: An unpredictable age” (2013)

Beard, “Some aspects of theories of mortality, cause of death analysis, forecasting and stochastic processes”, (1971)

Bennett et al “The future of life expectancy and life expectancy inequalities in England and Wales: Bayesian spatiotemporal forecasting”, (2015)

Bourbeau and Desjardins, “Mortality at extreme ages and data quality: The Canadian experience”, (2006)

Bronikowski et al, “Aging in the natural world”, (2011)

Cairns, Blake, Dowd & Kessler, “Phantoms Never Die: Living with unreliable mortality data”, (2014)

Carey et al, “Increased Risk of Acute Cardiovascular Events After Partner Bereavement: A Matched Cohort Study”, (2013)

Carriere and Valdez, “Annuity Valuation with Dependent Mortality”, (1996)

CMI Report 10 (1990)

CMI Report 16 (1998)

CMI Working Paper 21, “The graduation of the CMI 1999-2002 mortality experience; final “00” series mortality tables – assured lives”, (2006)

CMI Working Paper 22, “The graduation of the CMI 1999-2002 mortality experience; final “00” series mortality tables – annuitants and pensioners”, (2006)

CMI Working Paper 35, “The graduations of the CMI self-administered pension schemes 2000-2006 mortality experience: Final “S1” series of mortality tables”, (2008)

CMI Working Paper 42, “CMI life office mortality experience 2003-2006”, (2009)

CMI Working Paper 65, “Analysis of the mortality experience of pensioners of self-administered pension schemes for the period 2004-2011”, (2013)

CMI Working Paper 71, “Graduations of the CMI SAPS 2004-2011 mortality experience based on data collected by 30 Jun 2012: Final “S2” Series of Mortality Tables”, (2014)

CMI Working Paper 76, “Analysis of the mortality experience of pensioners of self-administered pension schemes for the period 2006-2013 based on data collected by 30 June 2014”, (2014)

CMI Working Paper 77, “Report of the Graduation and Modelling Working Party”, (2015)

CMI Working Paper 81, “Final “08” Series annuities tables”, (2015)

Currie, “Forecasting mortality at high ages” (2010)

Gallop, “Mortality at advanced ages in the United Kingdom”, (2002)

Gampe, “Human mortality beyond age 110”, (2010)

Gavrilov and Gavrilova, “The Biology of Life Span: A Quantitative Approach” (1991)

Gavrilov and Gavrilova “Mortality measurement at advanced ages- a study of SSA death master file”, (2011) Gavrilov and Gavrilova, “New Developments in the Biodemography of Aging and Longevity” (2014)

Gavrilov and Gavrilova, “Biodemography of mortality at extreme old ages: A study of rodents and humans” (2014)

Greenwood and Irwin, “Human Biology: A Record of Research” (1939)

Hill, Preston and Rosenwaike, “Age reporting among white Americans age 85+: Results of a record linkage study”, (2000)

Hoffmann, “Socioeconomic Differences in Old Age Mortality”, (2008)

Horiuchi and Wilmoth, “Deceleration in the age pattern of mortality at older ages”, (1998)

Kestenbaum and Ferguson “Supercentenarians in the United States”, (2010)

Kohler et al, “Comparative mortality levels among selected species of captive animal”, (2006)

Liang et al “Old age mortality in Japan: Does the socio-economic gradient interact with gender and age” (2002)

National Diabetes Audit 2011-2012, “Report 2: Complications and Mortality” (2013)

National End of Life Care Intelligence Network, “Deaths in older adults in England”, (2010)

ONS, “English Life Tables No.14, 1980-1982”, (1987)

ONS, “Decennial Life Tables – English Life Tables No.15, 1990-1992”, (1997)

ONS, “Decennial Life Tables – English Life Tables No.16, 2000-02”, (2009)

ONS, “National Population Projections, 2012-based Statistical Bulletin”, (2013)

ONS, “National Life Tables 2011-13, United Kingdom”, (2014)

ONS, “English Life Tables No.17, 2010 to 2012”, (2015)

Preston, Elo and Stewart, Effects of age misreporting on mortality estimates at older ages (1997)

Robine, Gampe and Vaupel, “IDL, the International Database on Longevity” (2005)

Stauffer “Simple tools for forecasts of population ageing in developed countries based on extrapolations of human mortality, fertility and migration”, (2002)

Thatcher, “The demography of centenarians in England and Wales”, (1999)

Thatcher et al, “The force of mortality at ages 80 to 120”, (1998)

Vaupel, “The Impact of Heterogeneity in Individual Frailty on the Dynamics of Mortality”, (1979)

Wilmoth, “Are mortality rates falling at extremely high ages? An investigation based on a model proposed by Coale & Kisker”, (1995)

Appendix A Terms of Reference

The CMI High Age Mortality Working Party was established in May 2014 with the Terms of Reference listed below.

1. **Context setting**
   - Broad indication of potential financial impact of misestimation of high age mortality
   - Data volumes required for credible assessments of high age mortality.

2. **Investigate and summarise published research on high age mortality**, including
   - The report of the CMI Graduation and Modelling Working Party
   - Commenting on likely absolute mortality rates at very high ages (compared with, for example, μ_{120} = 1 for the "S1"/"S2" Series tables) and the uncertainty around these estimates
   - Commenting on data and modelling issues particular to high age mortality data, for example: IBNR, data quality
   - Whether there is convergence with age (i.e. rates start to plateau)
   - Variation by sex, locale and socio-economic type (if any)
   - How high age mortality might vary in the future (noting the run-off to zero rates of change at high ages in the current CMI Mortality Projections Model)
   - Whether variation by socio-economic type at earlier ages adds materially to the uncertainty surrounding high age mortality extrapolation
   - Possible data sources for further investigation, with particular (but not exclusive) attention to the UK.

3. **In relation to existing data sources used by the CMI (SAPS, insurers, ONS)**
   - Issues with existing data sources specific to high age mortality, for example: ONS high age mortality methodology; effectiveness/prevalence of existence checking methods (current and past)
   - Determine what data the CMI could request from them that would assist with understanding high age mortality
   - Outline any selection criteria or other tactics the CMI might apply in order to reduce the scale of the data collection problem
   - Outline specific analysis that could be undertaken given that data.

4. **New research / analysis**

If/once the CMI has the data from 3 above, carry out the proposed analysis or such other analysis as appears appropriate with the objective of providing general guidance to CMI committees on how to incorporate high age mortality in their graduations and models.

Appendix B Graduations and extensions to the oldest ages

A wide variety of approaches have been used in the past for both graduating mortality rates and extending these to the oldest ages where the data are more scarce or less reliable. Three general approaches for extending graduations to the oldest ages are discussed in CMI Working Paper 77, namely:

- Using the same model as was used for the graduated age range. Some models will produce better results than others
- Blending the old age mortality rates into those from another known table. This would depend on a suitable known table being available.
- Applying a 'bolt-on' approach using a different formula for the older ages. This will typically be made to satisfy constraints of continuity and smoothness at the join and also to provide trajectories at the oldest ages which are thought plausible (e.g. that male mortality rates will be higher than female mortality rates for the same population type at all ages).

Extrapolation techniques

Thatcher, Kannisto and Vaupel (1998) provides a very good introduction to a number of alternative models for high-age mortality. The authors consider 7 models, which they describe as:

- Gompertz: $\mu_x = ae^{bx}$
- Makeham: $\mu_x = c + ae^{bx}$
- Logistic: $\mu_x = c + ae^{bx}/(1+ \alpha e^{bx})$
- Kannisto: $\mu_x = c + ae^{bx}/(1+ ae^{bx})$ *or* $logit(q_x)=ln(a)+bx$
- Weibull: $\mu_x = ax^b$
- Heligman & Pollard: $q_x = ae^{bx}/(1+ ae^{bx})$ *or* $logit(q_x)=ln(a)+bx$
- Quadratic: $ln(\mu_x) = a + bx + cx^2$

Excluding the Makeham model, as it is broadly the same as a Gompertz model at high ages the above six models were fitted using a maximum likelihood method to data for ages 80 to 98 collected from 13 countries (western European plus Japan) for the period 1960 to 1990. The resulting rates for ages 99 to 109 were compared against the observed values. The authors conclude that the Gompertz model gave estimates which were furthest away from the observed values followed by the Weibull model and then the Heligman & Pollard model. The other three models gave results far closer to the observed values.

Whilst the quadratic method is pragmatic it has little theoretical support and will usually fail at some ages; as the coefficient of the x^2 term will typically be negative, due to mortality deceleration (as described in section 4.1) resulting in mortality rates decreasing above some age.

The logistic model is more general than the other models and will work in situations where the others do not fit the data sufficiently well. It also has some explanatory backing; for example, it results from assuming that individuals are subject to the force of mortality in the form of Makeham's law but with the parameter *a* varying between individuals such that they have a gamma distribution at birth.

The Kannisto model is a special case of the logistic model which works if the constant *c* in the logistic model is small, the function μ_x can be fitted reasonably well by a function which tends to 1 and that the sex ratio of mortality rates is close to 1 at high ages. If these conditions hold it is easy to fit as it has one fewer parameter than the logistic model.

The paper concludes that the best models are the logistic model and its Kannisto approximation. On the data used for the analysis the fitted logistic model gave an estimated value of μ_{120} of between 0.7 and 1.0.

The favoured models all assume mortality decline (i.e. mortality rates increase more slowly than log-linearly at the highest ages). The element of shape by age is a recurring consideration in papers looking at high age mortality. However, not all authors agree that there is mortality decline at the oldest ages.

It should also be remembered that the results will be sensitive to the data and that results which held in 1998 may no longer do so (and some which performed less well may perform better) for more recent data. In particular, the Kannisto method assumes that μ_x tends to 1 (and hence q_x tends to around 63%). This may be a pragmatic solution that works with certain datasets and at certain times, but should not be expected to hold in general.

High age extensions in graduated tables

As referred to in section 4, the table below summarises the high age extension approaches used in a variety of UK and North American tables.

Table	Description	Comments	Summary of high age approach / constraints
ELT17M	Based on population data from England and Wales, ELT17, ELT16, ELT15 and ELT14 reflect experience over the periods 2010-12, 2000-02, 1990-92 and 1980-82 respectively.	A set of models are derived depending on a choice of threshold age, x_0, where the model transitions to the assumed model for old age mortality with a choice of a Gompertz or logistic model for the old age mortality. For both models, the graduated mortality rates for $x < x_0$ are provided by a generalised additive model. Each model is then evaluated on the basis of how well it predicts a set of data (the validation data) using a separate set of data (the training data). In this case the number of deaths by age for 2010 and 2012 formed the training data and the numbers of deaths for 2011 the validation data. Probabilities are then computed, which account for the extent to which the data support different values of threshold age, x_0, and each of the two models, using a Bayesian approach. The final graduations of m_x then represent a weighted average of the values produced by each model and threshold age. This method is carried out separately for males and for females.	Extension based on weighted average of values produced by Gompertz and logistic models with a range of transition ages. Probabilities for each model and transition age computed using a Bayesian approach based on actual data.
ELT16M		Uses data covering ages up to 108. Graduated mortality rates are determined using a variable-knot spline regression approach. There is no separate high age extension methodology; high age values are determined directly from the spline regression (with a high age constraint of m_{120}=2). With the exception of age 103 (where the crude q_x looks surprisingly low) the fitted rates adhere closely to the crude data.	**Continuation of graduation (with constraints)** • $m_{120} = 2$
ELT15M		Used data covering ages up to 103. Was extended above age 103 by extrapolating a linear regression of $ln(q_x/(1-q_x))$ fitted from age 85 to 103.	**Extension using logistic regression above age 103** • Extrapolation of a linear regression of $ln(q_x/(1-q_x))$ fitted from age 85 to 103
ELT14M		Used data covering ages up to 99. The crude q_x at ages 95 and above appear surprisingly low – in that they do not follow the usual pattern of q_x increasing with age – which is suggestive of data distortions at higher ages. As a result was extended above age 92 using a cubic polynomial fitted around the join and a constraint of $m_{105}=0.75$.	**Extension using cubic polynomial above age 92** • value, first and second derivatives at age 92 equal to those of quadratic defined by m_{90}, m_{91} and m_{92} • $m_{105} = 0.75$

Table	Description	Comments	Summary of high age approach / constraints
CPM2014M	Published by the Canadian Institute of Actuaries, these tables are based on Canadian registered pension plan experience over the period 1999 to 2008, projected to 2014. The tables are amounts based.	Graduated using amounts data up to age 100. The graduated table is markedly lighter than the crude mortality rates; this is due to the data (centred around early 2004) being projected to 2014 before being graduated. The 95% confidence intervals around the crude q_x increase markedly above age 95 (reflecting the lower volumes of data at those ages) and decrease in value at age 100, consistent with data distortions at higher ages (although as data is only published at sample ages this could just be statistical volatility). From age 103 to 114 the high age extension reflects a graduation of the broader Canadian population; at ages 95 to 102 a quartic polynomial is used to blend from the graduated rates to the high age rates.	**Extension using quartic polynomial for ages 95 to 102 and a graduation of population data from age 103 to 114** • Quartic polynomial constrained to match graduated rates at ages 92, 93 and 94 and graduated population data rates at ages 103 and 104 • Canadian population data graduated from age 95 to 110 using Whittaker-Henderson approach.
RPH-2014M	These tables are published by the Society of Actuaries' Retirement Plans Experience Committee and are based on uninsured private pension plans in the United States. The RP-2000 and RP-2014 tables reflect experience over the periods 1990-94 (projected to 2000) and 2004-08 (projected to 2014) respectively. The tables are primarily amounts based, although in the case of the RP-2014 series a headcount based variant is also produced and that has been used for the comparisons in Section 4.	Both RPH-2014M (the headcount based variant of the RP-2014M table) and RP-2000M are graduated using data up to age 100. The RP-2014M graduated table is lighter than the crude mortality rates; this is due to the graduated rates incorporating a projection to 2014 (whereas the data is centred around 2006). This feature is not apparent in the RP-2000M table, despite the rates being projected from 1992 (the central year of the underlying data) to 2000. In both cases high age extensions provides q_x values from ages 101 to 120; RPH-2014M uses a Kannisto regression, capped at a q_x of 0.5 and with Lagrange smoothing to blend into the graduation and RP-2000M uses a cubic then level extension designed to reach a q_x of 0.4 with a slope of 0 at the join. Both extensions provide q_x values appreciably lower than ELT16M.	**Extension using Kannisto regression above age 100, capped at a q_x of 0.5** • Kannisto regression (of crude μ_x) fitted over age range 75 to 104, using weighted nonlinear least squares procedure • q_x capped at 0.5 • Lagrange interpolation used to transition smoothly from graduated rates to extension
RP-2000M			**Extension using Cubic polynomial above age 100** • q_x matches graduated rates at $x=99$ and $x=100$ • slope of q_x is 0 at value of x where q_x reaches 0.4 and no inflection points between age 100 and that value of x • q_x set equal to 0.4 at ages above which it first attains 0.4
S2PML	Published by the CMI, these tables are based on data from UK occupational pension schemes. While a wide range of tables are provided, we focus in Section 4 on the lives-based all-	Both S2PML and S1PML are graduated using data up to age 95. The 95% confidence intervals around the crude q_x increase markedly above age 95 (reflecting the lower volumes of data at those ages). In addition the crude q_x decrease in value at ages around 100 and above (after levelling off in the high 90s), suggesting of data distortions at higher ages.	**Extension using cubic spline above age 95** • μ_x is continuous where the fitted curve meets its extension • first differential of μ_x is continuous at the same point, to ensure relatively smooth progression • $\mu_{120} = 1$

<table>
<tr><th>Table</th><th>Description</th><th>Comments</th><th>Summary of high age approach / constraints</th></tr>
<tr><td></td><td rowspan="2">pensioner (excluding dependants) tables.</td><td rowspan="2">In both cases a high age extension provides q_x values at ages 96 to 120. This extension provides q_x values appreciably lower than ELT16M.</td><td>• first differential of μ_x with respect to t (where $t = (x\text{-}70)/50$) is constrained to be 1 at age 120
(Certain other steps were taken to avoid anomalies between related tables – principally constraining the values of certain tables between 90 and 95 so they converged by 95, and constraining the value of other tables above age 95 to avoid crossings.)</td></tr>
<tr><td>S1PML</td><td>Extension using cubic spline above age 95
• μ_x is continuous at age 95
• first differential of μ_x is continuous at 95
• $\mu_{120} = 1$
• gradient (with respect to t where $t = (x\text{-}70)/50$) of the curve of μ_x is 1 at age 120</td></tr>
<tr><td>PML08</td><td rowspan="3">Also published by the CMI, these tables are based on data from UK life offices over the periods 1979-82 (“80” Series), 1991-94 (“92” Series), 1999-2002 (“00” Series) and 2007-2010 (“08” Series). Again a wide range of tables are provided; we focus in Section 4 on the lives-based combined-pensioner tables.</td><td rowspan="3">PML08 and PCML00 are graduated using data up to age 100 and 97 respectively. In contrast PML92 and PML80 were graduated using all available data (i.e. up to ages 108 and 107 respectively). As for the SAPS data, the 95% confidence intervals around the crude q_x increase markedly above age 95 (reflecting the lower volumes of data at those ages) and decrease in value at ages around 100 and above (after levelling off in the high 90s), suggesting of data distortions at higher ages.

For PML08 and PCML00, high age extensions provide q_x values at ages 90 to 120 and 98 to 120 respectively. High age values for PML92 are in contrast taken from the graduated curve (noting that the graduation parameters were adjusted somewhat to achieve consistency between males / females and lives / amounts). For PML80, q_x values at ages 91 and above effectively inherit their</td><td>Extension using non-linear interpolation above age 90
Same approach as for PCML00 (see below), but with $a=90$ and $c=0.95$ for PML08</td></tr>
<tr><td>PCML00</td><td>Extension using non-linear interpolation above age 97
At the oldest ages, values of μ_x for $x > a$ were blended into an arbitrary μ_{120} equal to 1 using the formula:
$$\mu_x = \frac{(120-x)^c}{(120-a)^c} \times \mu_a + \left(1 - \frac{(120-x)^c}{(120-a)^c}\right) \times \mu_{120}$$
For PCML00 $a=97$ and $c=0.80$</td></tr>
<tr><td>PML92</td><td>Continuation of graduation
Graduation parameters adjusted somewhat to achieve consistency between M/F and L/A tables</td></tr>
</table>

Table	Description	Comments	Summary of high age approach / constraints
PML80		shape from the corresponding amounts table, which in turn was adjusted at ages 93 and above to provide greater consistency with permanent assurances duration 2+ at the highest ages. The high age q_x values are (particularly for PML08, PCML00 and PML92) appreciably lower than those for ELT16M.	**Extension based on PMA80 above age 91** • μ_x set to be constant multiple of PMA80 table for $x>=91$ • multiple set so $\mu_{91}^{extended} = \mu_{91}^{graduated}$ (PMA80 table was itself adjusted at ages 93+ to provide greater consistency with permanent assurances duration 2+ at the highest ages)

Further detail on the high age extension methods underlying the ELT tables can be found in Gallop's 2002 paper, 'Mortality at high ages in the United Kingdom'. Details of the approaches underlying the CMI SAPS and life office tables can be found various CMI working papers (in particular CMI Working Paper 35 and CMI Working Paper 71 for SAPS tables, and CMI Report 10, CMI Report 16, CMI Working Paper 22 and CMI Working Paper 81 for life office tables). Details on the Canadian Institute of Actuaries' and Society of Actuaries' tables can be found via www.cia-ica.ca and www.soa.org respectively.

Appendix C Results from different datasets on the question of convergence of mortality with age

In this Appendix we summarise results from different datasets focusing on whether or not we observe convergence of mortality with increasing age. This has relevance to those who are considering fitting to "population" mortality at the oldest ages, as discussed in section 8.

Most research studies show converging mortality with increasing age. We present the evidence for convergence for lives with different (1) socio-economic circumstance or levels of deprivation, (2) pension size, (3) health status at retirement, (4) lifestyle choices, and (5) education. Some studies have shown no convergence and even divergence though these are in the minority (a number of studies are highlighted by Hoffman (2008)).

C.1 Social deprivation and socioeconomic circumstance

The National End of Life Care Intelligence Network is a government programme run through the Department of Health. The patients studied are close to death. They use the Income Deprivation Affecting Older People Index (IDAOPI). For England 2006-2008 we see little evidence of convergence across deprivation quintiles with increasing age though there is some crossover with the most deprived group, as shown in Figure C.1. A caveat on this study is that over one fifth of deaths over age 75 were in care homes (likely significantly higher for older age groups) which are more likely to be represented in more deprived areas, so uneven flows between deaths and exposures into care homes could distort results. Bale and Lu (2013) report socio-economic convergence is observed in larger datasets, including the General Practice Research Database and Office for National Statistics total English population.

Figure C.1: Comparison of deaths by deprivation quintile

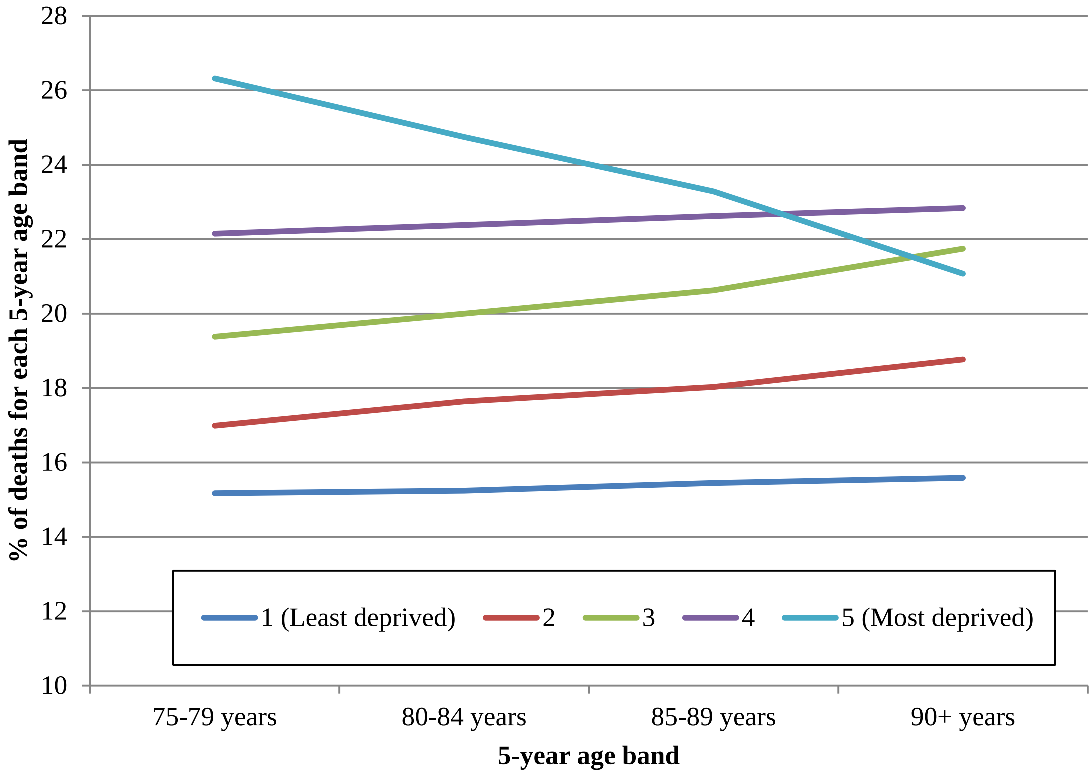

Source: Public Health England, 2006-2008. Information request, from National end of life care intelligence network, "Deaths in Older Adults in England", (2010).

C.2 Pension Size

Convergence by pension size has been observed in SAPS investigations, as can be seen from the executive summary accompanying Working Paper 76, covering the period from 1 January 2006 to 31 December 2013. Figures C.2 and C.3 show that, as age increases, the difference in mortality between the pension bands at higher ages narrows. The differences between pension bands are less pronounced for females and the convergence is less distinct.

Figure C.2: 100 A/E values for Male Pensioners Amounts compared to S2PMA

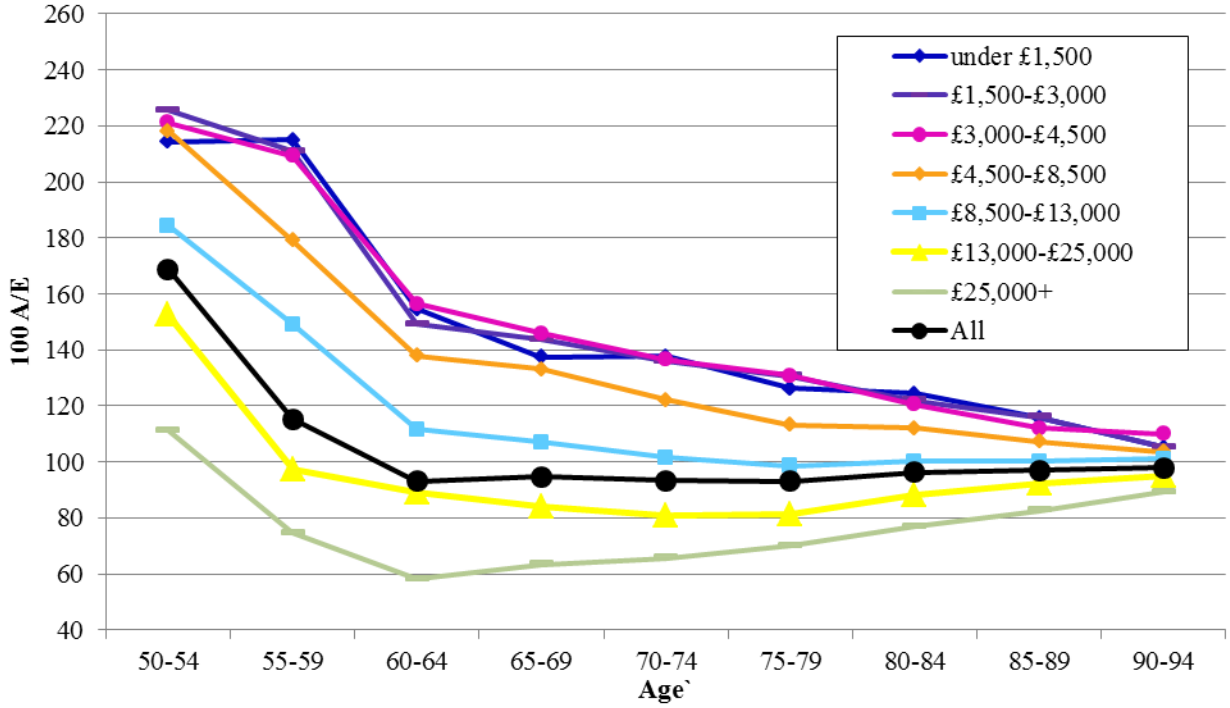

Source: CMI Working Paper 76, Executive Summary

Figure C.3: 100 A/E values for Female Pensioners Amounts compared to S2PFA

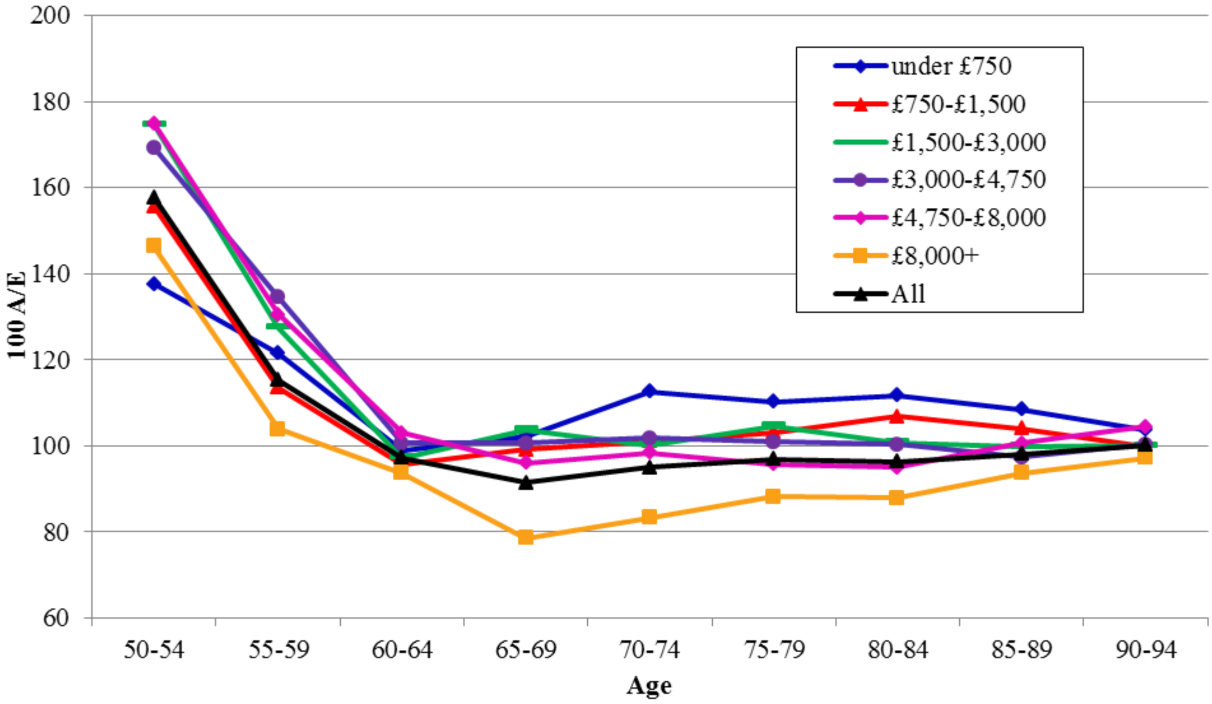

Source: CMI Working Paper 76, Executive Summary

C.3 Health status at retirement

There is a view that ill-health mortality gradually converges to that of standard health. This view is often drawn from experience analyses which appear to show such a convergence with increasing age, as seen in Figures C.4 and C.5, based on the dataset used for the graduation of the SAPS "S2" series tables. There are two possible explanations for this outcome: (1) ill-health retirees may recover from their ailment (2) normal health retirees fall ill as they age. The message is not that ill-health wears off with time; within the ill-health retirees we would expect the lightly impaired lives to be the survivors.

Figure C.4: Death rate in normal and ill heath males in the SAPS "S2" Series dataset

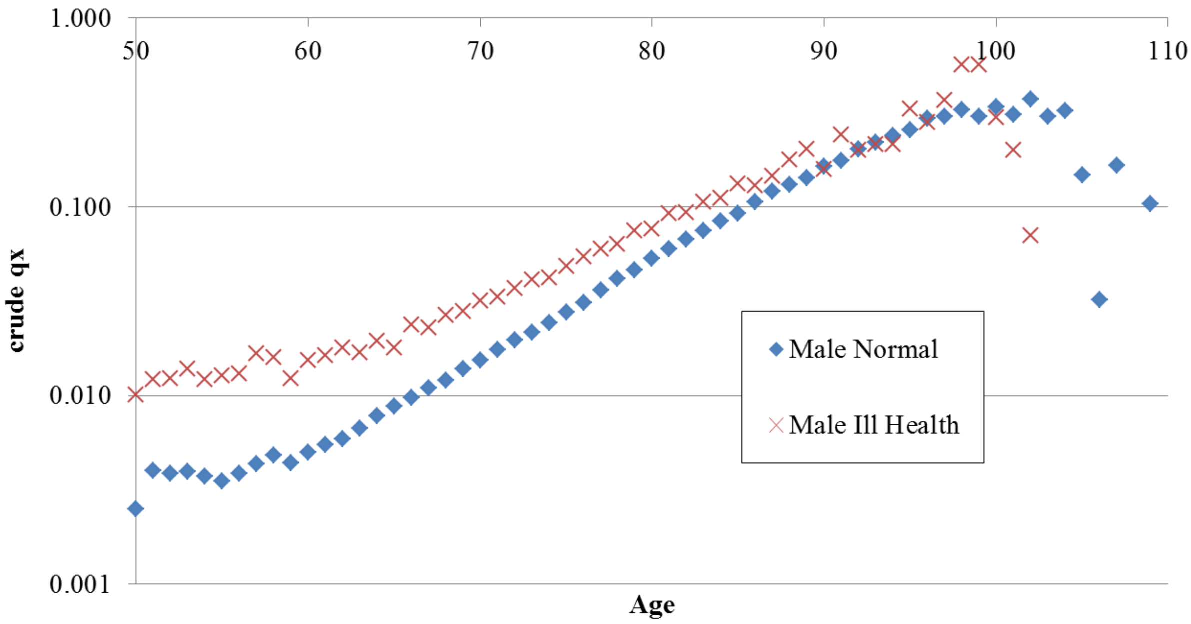

Figure C.5: Death rate in normal and ill heath females in the SAPS "S2" Series dataset

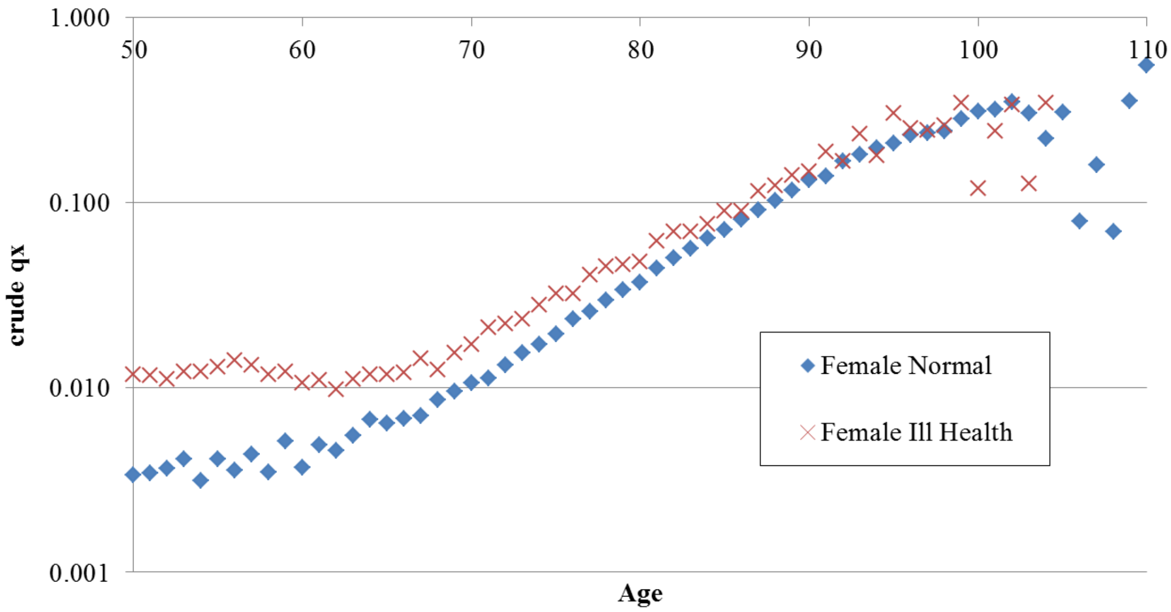

Definition of ill-health retirement

It is possible that ill health retirement rates have fallen since the 1970s as scheme practices and rules have tightened up the eligibility criteria for granting such a benefit. One might question whether some of the "ill-health" retirements of the 70's and 80's were actually in normal health. If lives were mislabeled, or re-defined over a period of time then this might result in the ill-health experience converging faster to the normal experience at the older ages of an experience analysis.

Progressive diseases

A progressive disease is one whose course in most cases is the worsening, growth, or spread of the disease. The means the mortality for progressive diseases increases with duration from diagnosis (e.g. diabetes, Alzheimer's). We investigate if there is a convergence in relative risk even for progressive diseases.

As an example of how relative risk ratio develops with age for a progressive disease, we have reproduced a graph of the relative risk ratio for type 1 and type 2 diabetes sourced from the National Diabetes Audit 2011-2012 - Report 2: Complications and Mortality. This shows that for type 2 diabetes (which is more common) the relative risk is around 1.7 for ages 35-64, reducing to 1.4 for ages 65-74, and is around 1.1 for the population above age 85. It is a matter of opinion whether it gets to 1 at an older age.

We have not investigated how relative risk behaves for other progressive diseases so the convergence we see for diabetes may not necessarily apply.

Figure C.6: Age-specific mortality rate ratios by type of diabetes and sex

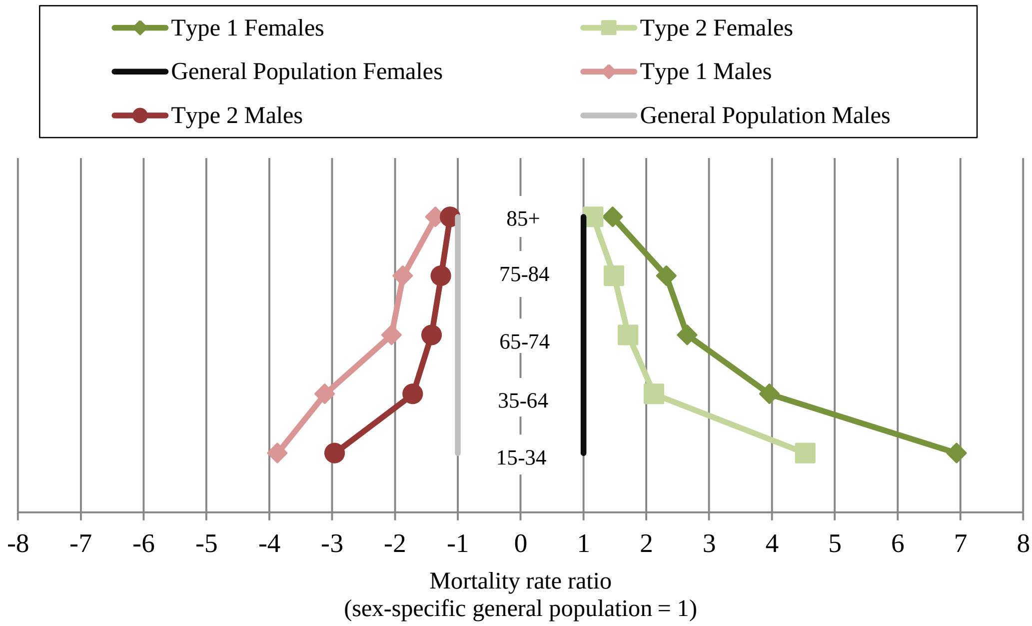

Source: National Diabetes Audit 2011-2012 Report 2: Complications and Mortality

Note that if one changes the focus from relative mortality to excess mortality the picture could be quite different. Received (but unproven) medical wisdom is that organ damage accumulates with duration since diabetes diagnosis. Hence in general excess mortality for a diabetic increases with advancing age but relative mortality reduces as the excess does not increase as quickly as average all-cause mortality.

C.4 Lifestyle choices

Data on the relative mortality between smokers and non-smokers was collected for CMI Working Paper 42. The excess mortality index for smokers versus non-smokers for durations 2+ decreases after the 71-75 age group (apart from a very small increase at ages 86-90 which could be due to low levels of data), as seen in Figure C.7, which supports the idea that excess mortality converges.

Note: the Excess Mortality Index is calculated as the percentage by which the standardised mortality ratio for smokers exceeds the corresponding ratio for non-smokers, i.e.:

100 x (100A/E Smokers / 100A/E Non-smokers)

Figure C.7: Excess mortality amongst male smokers compared to male non-smokers

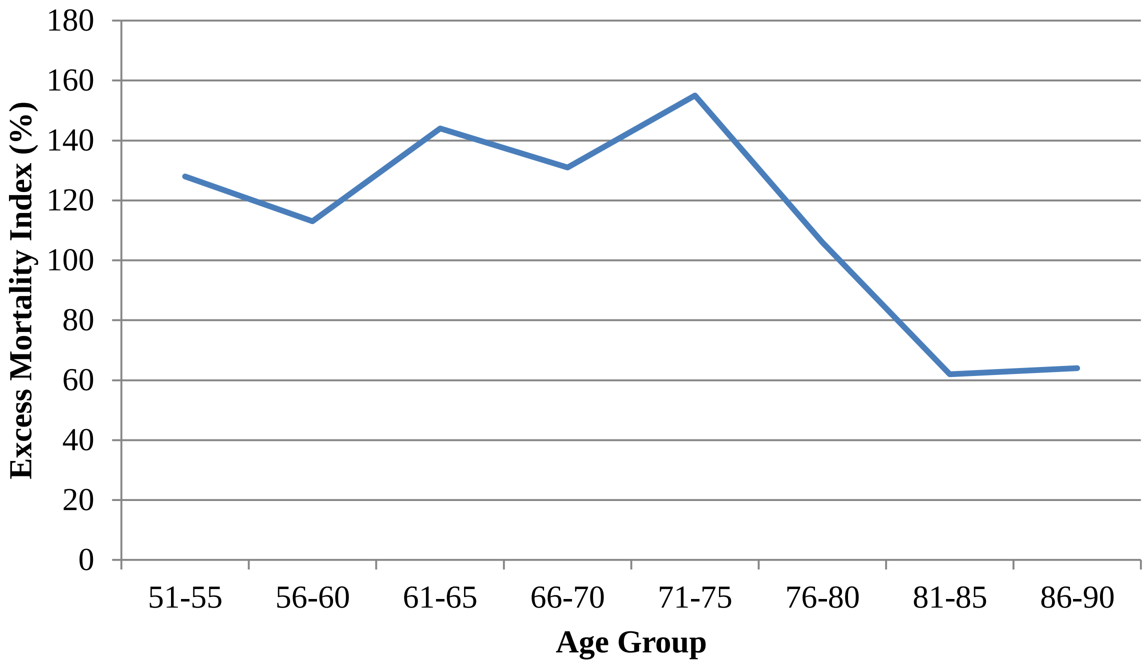

Source: CMI Working Paper 42.

C.5 Education

Liang et al (2002) analysed mortality differences over age with a panel study of health and wellbeing of older adults (60+) in Japan with five waves from 1987 to 1999. Their main finding is an educational mortality crossover for men. That is to say that, at advanced ages, those men with less education live longer than those with higher education, whereas this is reversed at younger ages and a gap maintained for females.

Contrary to this at the 2014 Longevity and Mortality Symposium plenary session Robert L Brown presented the results of a US study which does not appear to show convergence by education. If fewer people went to university or college say in the 1930's than in 1970's then this might explain why higher educated lives have lighter mortality far into the advanced ages as the former is expected to be a more select group. Indeed, it is not clear whether the data has been standardised by other factors to isolate the "education" effect.