ISSN 2044-3145

Continuous Mortality Investigation
Institute and Faculty of Actuaries

High Age Mortality Working Party

WORKING PAPER 100

A second report on high age mortality

June 2017

Executive Summary

Background

Estimating the population at high ages in England & Wales (E&W) is important as a prerequisite for national population projections, estimating current mortality rates by age (and hence life expectancy) and estimating past patterns of mortality improvement. These have public policy implications for pension provision, social care and health care, and are of direct interest to pension providers and life insurance companies in the private sector.

The CMI has a particular interest in this field, as it feeds into:

- the calibration of the CMI Mortality Projections Model for estimating past (and future) improvements; and
- the determination of population mortality curves that may be used to close-off portfolio mortality tables at high ages, where the underlying datasets lack credibility, (as described in Section 4 of this paper), potentially including the Self-Administered Pension Schemes (SAPS) and Annuities base mortality tables.

In Working Paper 85, the High Age Mortality Working Party summarised published analysis on very high age mortality, as well as indicative modelling on extinct cohort mortality and the potential impact of age mis-statement and delays in death reporting. For clarification, when discussing high age mortality we generally mean mortality at around age 90 and above (though this may vary from case to case according to the nature of the data under consideration).

Population exposure modelling

The Working Party has reviewed the Kannisto-Thatcher (KT) method in the form currently used by the Office for National Statistics (ONS), and analysed variants designed to address a number of limitations in the current approach.

Our analysis suggests that there may be merit in adopting the following modifications to the current ONS methodology for estimating high age population exposures for E&W:

- Refine the projection of survivor ratios in the KT method, by allowing for local mortality trends over time and correspondingly reducing the number of birth-year cohorts in the survivor ratio from 5 to 2;
- Extend the high age population method down to a lower age than 90 (for example, age 85) to avoid placing undue reliance on the underlying census-based estimates in this age range;
- Incorporate a more sophisticated approach to adjusting the input deaths data to a 1 January age definition (as required by the KT method when using deaths data by age at death) using Lexis triangles; and
- Convert the resulting population estimates to exposures by smoothing the final estimates as a pragmatic solution to issues arising from uneven birth and death distributions during the year.

If all of these changes were adopted, we estimate that male period life expectancy based on ONS E&W data would be reduced at the highest ages. For females, the reduction would be smaller (with modest increases in life expectancy at some ages, depending on the approach to graduation). The impact on life expectancy at age 85 using ungraduated mortality rates is a reduction of 0.1% for males and an increase of 0.1% for females based on ONS E&W data over the period 2011 to 2015. The impact on life expectancy is more material at age 95, where we observe a 2.1% reduction for males and a 1.1% reduction for females.

The conclusions of our research into population exposure modelling and selected testing results are included in this paper. The supplementary Technical Paper sets out more detail on the approach, justification and testing.

Mortality at very high ages

Working Paper 85 highlighted that there were differing views as to whether the progression of mortality rates into very high ages, i.e. those aged 100 and above, follows an exponential (Gompertz) progression or whether there is an element of deceleration in rates. To further this discussion, and to help users of mortality curves to set mortality rates at the highest ages, we have considered three further papers that have studied independent datasets of mortality rates at very high ages:

- Gavrilov and Gavrilova (2015) analysed mortality data from the International Database on Longevity (IDL) and concluded that a Gompertz curve is the best fit for high age mortality.
- Ouellette and Bourbeau (2014) studied the mortality of Canadian centenarians using Quebec's church parish registers. The authors were able to match birth dates on death certificates with baptismal certificates for the majority of cases studied which gave a high degree of certainty around the data quality. They concluded that there is observed deceleration in the death rates at older ages.
- Rau et al (2016) fitted Gompertz, Gompertz-Makeham, Gamma-Gompertz and Gamma-Gompertz-Makeham models to male and female data from seven countries and across different time periods. The authors concluded that:
  - models with mortality deceleration provided the best fit across virtually all time periods and countries; and
  - the level of the plateau for μ_x was around 0.8 for females and 1.2 for males although the male results showed significant variation.

The Working Party has analysed the log of mortality rates determined independently from the three approaches. We conclude that it is hard to justify a log-linear fit to the E&W data in light of the data sources we have collected and the confidence intervals around them. Based on the papers we have reviewed and analysis we have performed, we believe that a mortality curve which makes allowance for deceleration at advanced ages is appropriate for period mortality, i.e. mortality rates over the period analysed with no allowance for future improvements, at a population level. We also believe that currently a value of μ_{120} of 1 is justifiable and a reasonable assumption.

Closing-off mortality tables

The Working Party has determined a framework with which to close-off mortality tables at high ages, i.e. a framework for extending graduated rates to higher ages, where the data lacks credibility, and to determine the mortality rate at the highest age for which the table provides mortality rates. This framework has been constructed to assist the mortality-focused CMI committees with setting mortality rates for future tables.

The Working Party recommends the following approach:

- Graduated mortality rates are set principally by reference to the experience of the portfolio in question for all of the ages at which sufficient, credible data is available for graduation. As such the convergence from portfolio to population mortality can then be assessed.
- As age increases and portfolio data becomes insufficient (in volume and / or quality) to reliably graduate mortality rates, wider data (e.g. relevant population data) needs to be used to inform the graduation / extension. A view must therefore be formed on the degree to which there is convergence between the portfolio rates being graduated and the population data in this age range and the rate at which this occurs.
- Volumes of mortality data reduce rapidly at very high ages, and there reaches a certain age where even population data becomes insufficient to set the mortality rate curve. Expert judgement informed by any available data as well as plausible views as to the shape and level of mortality should be applied here. This approach will in turn imply an assumed ultimate age and mortality rate at the maximum age of the mortality table. The intended approach should seek to deliver consistent mortality rates produced by the various CMI mortality committees at very high ages.

This framework can be extended to determine extensions for a sub-portfolio relative to a main portfolio to ensure consistency in graduating rates and approach.

As part of this analysis the Working Party has performed case studies on recent SAPS and Annuities mortality data.

Next steps

The Working Party intends to analyse mortality data at high ages for members of very large UK pension schemes, and to report back on insights, subject to sufficient data volumes.

Contents

1. Introduction

Estimating the population at high ages in England and Wales (E&W) is important as a prerequisite for national population projections, estimating current mortality rates by age (and hence life expectancy) and past patterns of mortality improvement. These have public policy implications for pension provision, social care and health care, and are of direct interest to pension providers and life insurance companies in the private sector.

The CMI has a particular interest in this field, as it feeds into:

- the calibration of the CMI Mortality Projections Model for estimating past (and future) improvements; and
- the determination of population mortality curves that may be used to set portfolio rates for portfolio curves at high ages (as described in Section 4 of this paper), potentially including the SAPS and Annuities base mortality tables at high ages where the underlying datasets lack credibility.

In Working Paper 85, the High Age Mortality Working Party summarised published analysis on very high age mortality, as well as indicative modelling on extinct cohort mortality and the potential impact of age mis-statement and delays in death reporting. For clarification, when discussing high age mortality we generally mean mortality at around age 90 and above (though this may vary from case to case according to the nature of the data under consideration).

This paper, and the supplementary Technical Paper, build on the work performed in Working Paper 85 by exploring further the following areas:

- Population exposure modelling: In Section 2, we set out our proposed framework for modelling population exposures at high ages. We start by reviewing the approach currently used by the Office for National Statistics (ONS) and its limitations and then move on to summarise the results of testing alternative approaches designed to address these. In particular we explore further:
    - an explicit allowance for mortality trend in the Kannisto-Thatcher (KT) approach;
    - the impact of reducing the age above which the KT approach is deployed (the "join age");
    - adjustments to deaths data to allow for a more sophisticated "Lexis triangle" approach to determine age at 1 January death counts from age at death input data; and
    - exposure adjustments to smooth out abnormal population exposures.
- The shape of mortality at very high ages: In Section 3, we explore further the topic of mortality deceleration at high ages by analysing data relating to three international mortality studies.
- Closing-off mortality tables: In Section 4, we set out a proposed framework to setting mortality rates at very high ages, in particular exploring convergence between portfolio and population mortality.

Given the technical content of the population exposure modelling, we summarise our findings in this paper and provide more detailed methodology and analytics for this specific topic in the supplementary Technical Paper.

A full list of the previous research and other material reviewed in this paper can be found in the References section at the end of this paper.

We have focused on the mortality experience of pensions and annuities products, with limited consideration of the issues affecting whole of life products or other potential uses to the extent that these differ. Users should consider to what extent the features of this framework are suitable for protection-based products.

This paper has been prepared by the CMI High Age Mortality Working Party. The members of the Working Party involved in the preparation of this paper were Steve Bale (Chair), Carl Campbell, Mark Cooper, Andrew Gaches, Adrian Gallop, Andy Harding, Richard Lamb and Anny Sun. We would be very pleased to receive any comments or questions on this paper; these can be sent via e-mail to HighAgeMortality@cmilimited.co.uk.

The Working Party wishes to acknowledge and thank:

- Neil Robjohns, who helped shape the avenues of investigation performed by the Working Party and provided invaluable review and commentary on the research as it developed.

- Angele Storey and Pamela Cobb from the ONS Demographic Analysis Unit, who provided review and feedback on an early draft of the supplementary Technical Paper.
- Professor Andrew Cairns of Heriot-Watt University, who provided information on the distributional correction adjustments and Bayesian exposure smoothing framework proposed by Cairns et al in their ‘Phantoms never die’ paper (including running some of our test scenarios through the Bayesian framework) and reviewed an early draft of the supplementary Technical Paper.
- Leonid Gavrilov and Natalia Gavrilova who provided additional information on their modelling techniques and data, which helped us to understand their methodology as part of the investigation into the shape of mortality at very high ages.
- Piero Cocevar, Deborah Cooper, Tim Gordon and Steven Rimmer for their valuable review of draft versions of this paper and the supplementary Technical Paper.

2. Modelling population exposures

In this section, we set out our proposed framework for modelling population exposures at high ages. We start by reviewing the approach currently used by the ONS and its limitations and then move on to summarise the results of testing alternative approaches designed to address these.

More detail around the approach, justification and testing can be found in the supplementary Technical Paper.

2.1 Background and current approach

The ONS currently uses the KT method to estimate population numbers at the highest ages.

This method is often used where the official population data by single year of age is assumed to be unreliable at the highest ages but, in contrast, the reporting of deaths is assumed to be far more accurate. This can be considered the case in E&W as it is a legal requirement to register all deaths occurring, whereas census data is subject to a range of known recording and processing issues and in any case is gathered only at 10-year intervals.

The methodology used by the ONS to determine its high age population estimates for E&W is set out online. Chart 2.1 illustrates how the ONS's current implementation of the KT method works.

- For ages 89 and below, the census-based figures by single year of age are adopted without adjustment. These figures are mid-year estimates produced by rolling the decennial census counts forward allowing for ageing, births, deaths and migration (the cohort component method).
- For ages 90 and above, the KT method is applied using deaths data to estimate the single year of age population counts. Ideally, death occurrence data would always be used rather than death registration data, as the former records when deaths actually occurred, rather than when deaths were registered. However, in practice, deaths data here means registration data to 1992 (when occurrence data is not available) and occurrence data from 1993 onwards (except for the most recent year which is registration data, due to registration data being available more quickly, to allow an earlier release). In Table 2.1, we note that the ONS expect the difference between occurrence and registration data to be negligible in practice. The age from which the KT method is applied is known as the join age throughout the remainder of this paper.
- In the diagram, it is the gold area for which we need to estimate population exposures. The population estimates in the rest of the figure are known, or have already been determined. Lives aged above 90 in the white and the blue areas are believed to be fully extinct (i.e. all lives are assumed to have died) and so the deaths information can be used to calculate historical exposures.
- The current ONS method takes a parallelogram of past data (the blue area) comprising m birth cohorts, each with k years of prior deaths data, and calculates the average survivor ratio (survivors ÷ deaths) for point A based on that set of cohorts. The current ONS methodology uses 5 years of past deaths data over 5 years of birth cohorts, i.e. k = 5 and m = 5.
- It then applies the calculated survivor ratio to the single birth cohort containing point B, multiplying up by the known deaths count over the prior k ages in the cohort to give the population estimate for point B. Once this first exposure estimate has been calculated, population estimates at previous ages along the cohort can be reconstructed by adding back in the recorded deaths. Then the completed figures for that cohort can be used to estimate the population of the next cohort, and subsequent cohorts until exposure estimates for the whole of the gold triangle in Chart 2.1 have been calculated.
- Two further adjustments are made to ensure that the total population estimates for ages 90 and above calculated by the KT methodology are the same as the official population estimates for each year:
  - The total population estimate for ages 90+ in the final year is constrained to the census-based figure for that year. This implies a correction factor c (in the language of the original KT method), which represents the assumed constant year-on-year adjustment factor that applies to the survivor ratios when projecting them, and should ideally be close to 1. In this paper we

alternatively refer to the final year balancing adjustment, defined as c – 1. This will have a knock-on impact to estimates calculated using the survivor ratio approach.

- A uniform annual scaling to the resulting population estimates at ages 90+ in each previous year is then applied to constrain the 90+ population total to the census-based figure for that year.

Chart 2.1: Current ONS implementation of KT method

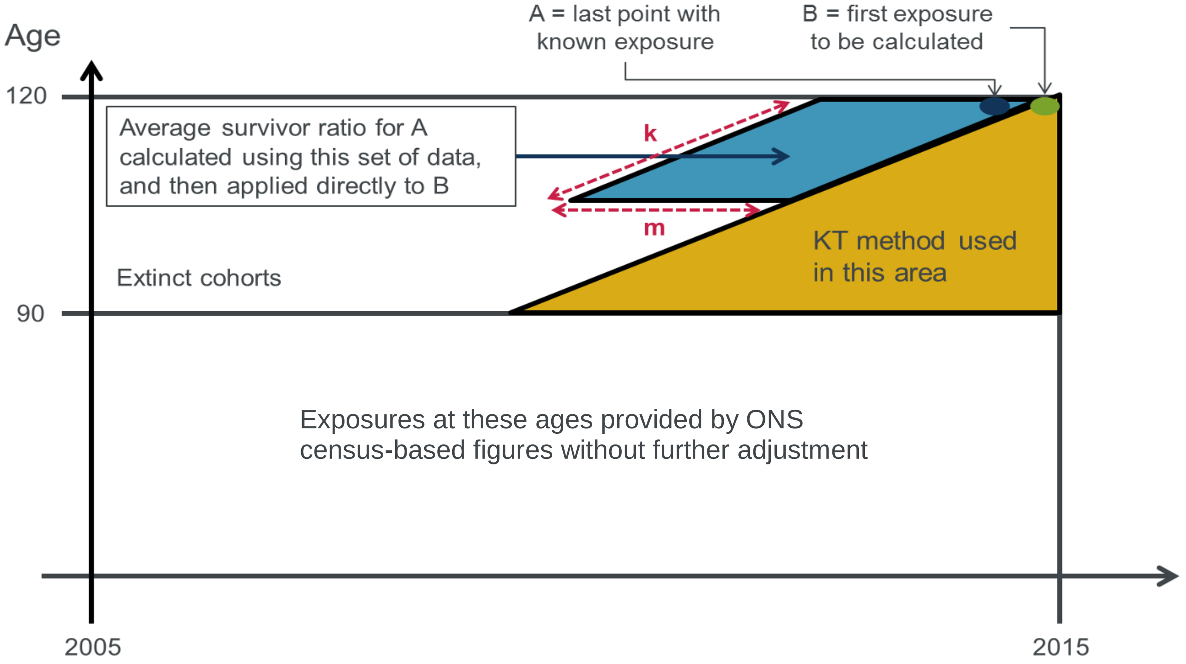

The population estimates over age 104 are then grouped into a single 105+ estimate for publication.

In December 2016, the ONS released their paper 'Accuracy of official high age population estimates in E&W an evaluation'. Their findings are summarised in Table 2.1 and we consider these are consistent both with the findings of our earlier work in Working Paper 85 and the work that is presented in this paper.

Table 2.1: Accuracy of official high age population estimates in E&W

Theme	Analysis and Implications
Death reporting	The ONS reported on an exercise by the University of Southampton to validate deaths data recording for a sample of deaths in E&W of semi-super-centenarians (aged 105 to 109) born within E&W and who died during the period 2000 to 2014. 1,141 deaths were included in the sample and of these 96% had a date of birth match between birth and death records, with a further 1-2% cases being partially validated with a small difference in dates. Whilst the sample set is small and restricted to deaths for ages 105 and above, it indicates a very high level of date quality validation for those born in E&W. Validation did not cover deaths of those born outside of E&W; around 13% of male deaths and 8% of female deaths aged 105 and over during 2000 to 2014 were of people born outside E&W. **The ultimate reporting of deaths is assumed to be complete**, given that it is a legal requirement for all deaths occurring in E&W to be registered. In addition, the ONS paper included analysis of the differences between annual occurrence and registration counts at high ages and concluded that both the numbers and distribution of deaths were very similar, with minimal impact on the KT estimates regardless of which type of data is used.
Impact of migration	The ONS has analysed the impact of both cross-border UK and international migration on population estimates. This is an important theme as extinct generation mortality modelling assumes that migration has a negligible impact. The work performed by the ONS concluded that levels of migration intra UK are very small for ages 90+. Cross-border UK migration from patient register data is of a similar magnitude to estimates from the 2011 Census, and is very small relative to the UK population data indicating that the quality of migration data at higher ages is acceptable. **The ONS has concluded that migration has a minimal impact on the mid-year population estimates of lives aged 90+.**
Kannisto-Thatcher modelling	The KT method relies on good quality deaths data and assumes negligible migration to determine population estimates at very high ages. The KT method was tested using population mortality data from Sweden and Finland for the period 2002-14 and where the deaths data are available in the format required. The KT estimates were compared to population registers for each country for each individual age as at the start or end of each calendar year during this period. The difference between the census registers and the KT estimates was minor. The ONS has also assessed the correction factor that is used to adjust the estimates to ensure consistency with the total estimate for 90+, using scenario testing of synthetic populations. **It was found that the correction factor was close to 1 (i.e. no adjustment) when mortality did not change over time, but deviated by up to 15% in either direction under mortality trend scenarios**.
Comparison against administrative data	The ONS compared the KT estimates for E&W against a dataset created by combining administrative datasets sourced from the National Health Service, the Department for Work and Pensions and the Higher Education Statistics Agency. For ages 90+, the administrative dataset returns population estimates for 2011 that are broadly similar to the unconstrained KT estimates. **This suggests to us that the unconstrained KT estimates are already in relatively close agreement with independent population estimates; therefore minimal population constraints using the final year balancing adjustment and annual scaling may be preferable.**

Members of the Working Party were invited by the ONS to attend their workshop in early 2016 where initial results for this paper were shared and discussed. This workshop and subsequent discussions with the ONS have helped shape our further analysis in the following two key areas:

- Mortality trend: In recent years we have observed heavier and more volatile mortality experience in the E&W population than had been experienced in the early 2000s. An allowance for a trend in mortality in recent years has been explored.
- Join age: The ONS has not analysed the impact of assuming a different join age to the currently assumed age 90. We have tested the impact of assuming a range of join ages.

In addition, the Working Party has explored the following refinements to modelling population exposures at very high ages:

- Parameters k and m: We have tested alternative values of m (the number of birth cohorts) and k (the number of past ages in each cohort over which deaths are summed) in the survivor ratios, compared with the current values (k,m) = (5,5).
- Transformation of deaths data: We have adopted a more sophisticated approach to determining age at 1 January death counts from the age at death input data (using a Lexis triangle decomposition).
- Smoothing distributional anomalies: We have considered the impact of adjusting the modelled population exposures for convexity, uneven birth patterns and other distributional issues (or applied pragmatic smoothing to similar effect).

In order to compare the performance of different variants in our testing, we have applied a set of numerical tests designed to draw out the desirable features. Details of the diagnostics can be found in the supplementary Technical Paper, but in summary these are:

- **Impact diagnostics**: The first question is whether a variant actually affects the population exposures to a material extent, which we address by showing period life expectancy estimates by age and population estimates by age relative to the current ONS approach.
- **Performance diagnostics**: These act as a guide to help determine whether variants perform 'better' or 'worse' than one another:
  - *Consistency with official population estimates*: Where better performance is indicated by either the final year balancing adjustments or the average annual scaling adjustment being closer to zero. Note that this does not in itself guarantee superior predictive performance (since the constraining population total may itself be inaccurate), but it does provide some comfort that the method being adopted is compatible with the constraint applied.
  - *Internal consistency:* Several measures including smoothness of mortality rates across join age, a cohort inconsistency metric and Cairns Blake Dowd Kessler (CBDK) diagnostics.
  - *Predictive performance:* Testing performance against six synthetic populations, where the 'correct answer' is known, unlike actual population data where the population numbers are estimates.

Our testing uses both ONS data and synthetic populations, as appropriate. Details of the six synthetic populations used are set out in Table 2.2.

Table 2.2: Synthetic populations modelled

Synthetic population	Mortality assumption
1. Mortality reduction 2% p.a.	Mortality rates reduce by 2.0% per annum.
2. Mortality increase 2% p.a.	Mortality rates increase by 2.0% per annum.
3. 2015 shock	No change in age-specific mortality rates from year to year, except for a one-off shock in 2015 where rates are assumed to be 5.0% higher.
4. Mortality reduction 2% p.a. (with cohort effects)	Mortality rates reduce by 2.0% per annum, but with higher rates of improvement for those born in 1910 to 1921 inclusive. Also includes small period effects and a slowing down in the higher cohort improvements and reductions after age 94 to zero improvements at age 114. The motivation for locating these effects in cohorts born between 1910 and 1921 is to produce a meaningful impact on the 90+ population over the final years of the dataset. Incorporating features akin to the golden cohort (i.e. lives born between World War I and World War II) observed in the actual E&W population would not have achieved this effect.
5. Raw E&W	Mortality rates determined with reference to mortality improvements for E&W females from 1970 onwards.
6. Smoothed E&W	As for population 5, but based on *smoothed* improvement rates for E&W females (from the CMI Mortality Projections Model, CMI_2015) from 1991 onwards

2.2 Summary of results of testing variants

More details of the testing results can be found in the supplementary Technical Paper.

Trend and varying k and m

In their December 2016 paper, the ONS noted that there may be merit in adopting a more sophisticated approach, with explicit allowance for survivor ratio trends in the KT method projection. We have tested a version of the KT method which allows for a simple 5-year linear extrapolation of the trend in past survivor ratios for the final calendar year at each age. This is illustrated in Chart 2.2, where rather than using a single blue parallelogram to determine the survivor ratio, we now look at a sequence of blue parallelograms through time and extrapolate the (linear) trend in survivor ratios to point B.

Chart 2.2: KT method with trend

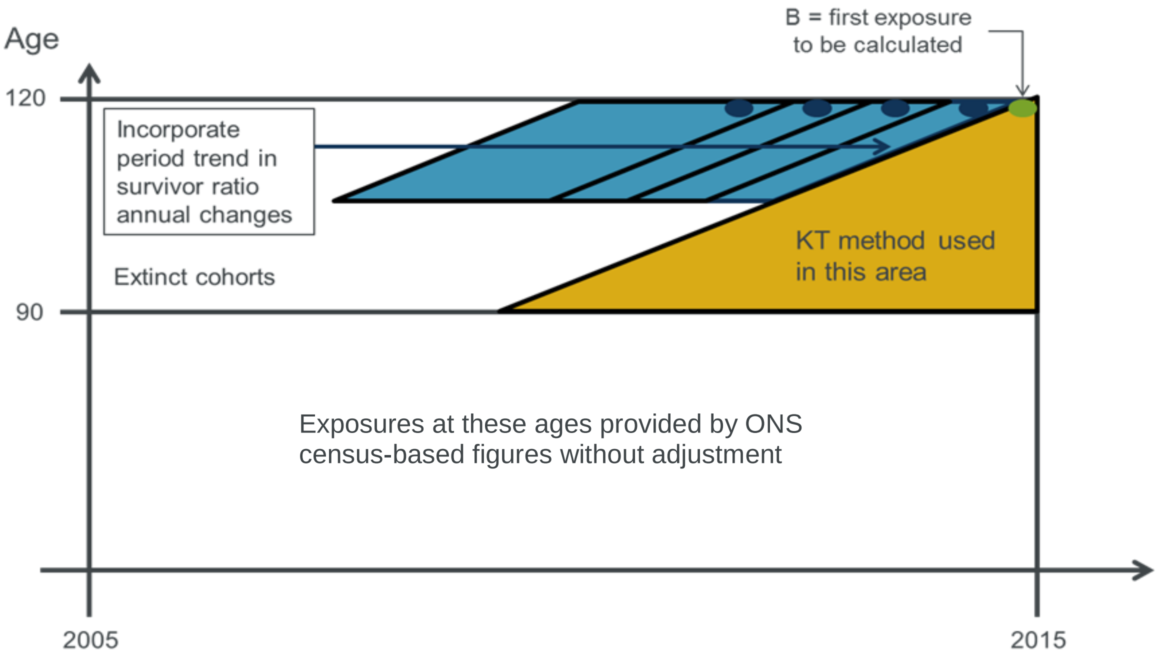

In addition to incorporating trend, we have also considered varying the values of m (the number of past cohorts) and k (the number of prior ages along each cohort over which the deaths are summed) in the historical survivor ratios. The current ONS method uses values of 5 for both k and m. In the 2002 paper by Thatcher et al, in which the KT method was first proposed, alternative values of k and m in the range 1 to 10 were considered and it was noted that there was ‘no clear optimal combination’. We have tested alternative parameterisations and found that, particularly following the introduction of mortality trend, a reduced value of m = 2 resulted in improved diagnostic tests. Chart 2.3 demonstrates one of these diagnostic tests, showing the impact of incorporating trend and then m = 2 on the final year balancing adjustment required for our six synthetic populations (where a final year balancing adjustment closer to 0 is considered an improvement). The adjustment is reduced overall in the majority of cases.

Chart 2.3: Final year balancing adjustment testing against synthetic populations

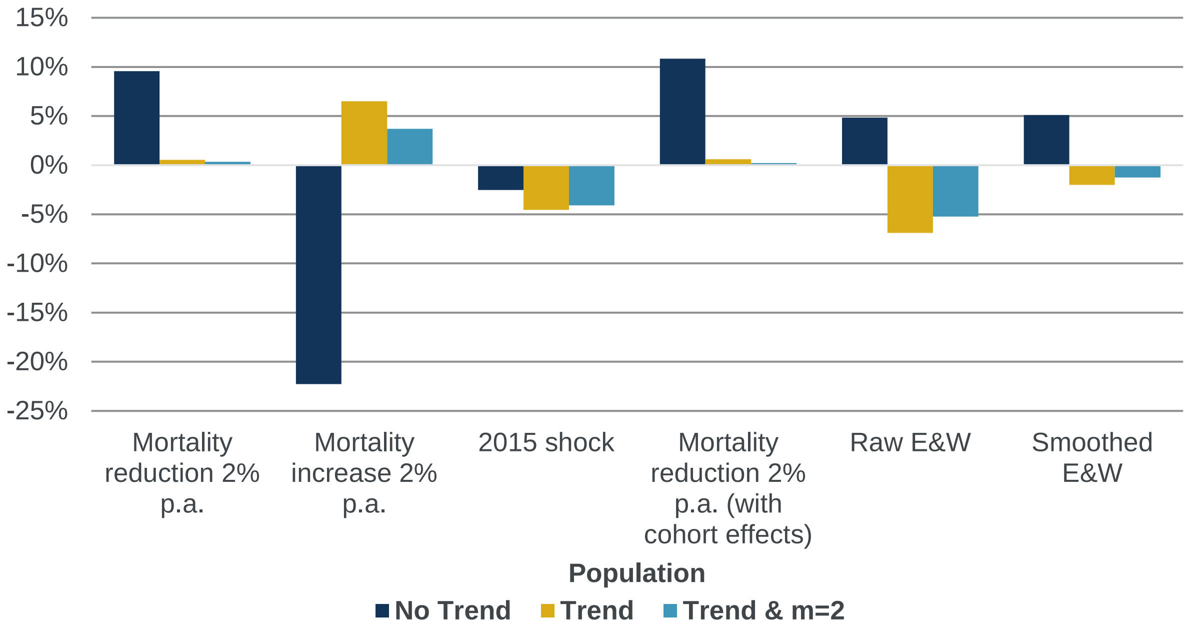

Join Age

The ONS currently applies the KT method down to age 90 and constrains the resulting estimates to match the census-based population total for ages 90+ in each year.

The rationale behind this approach is that the census-based figures are assumed to be reliable for each single year of age up to 89 and in aggregate for the age group 90+. However, there is an accumulation of evidence suggesting that these figures may not be accurate. For example, comparison with a range of alternative sources suggested that ‘there may be slightly too many people in the 2011 Census estimate at the oldest ages’. The ONS’s December 2016 paper also notes ‘discontinuities in population estimates at the age 89/90 boundary’ of the KT method, which may be caused by constraining the 90+ figures to an inaccurate total.

We would expect the reliability of census-based figures to improve at ages below 90, as the population counts are larger and less susceptible to some of the specific census recording issues identified by the ONS, but reducing the join age too far may introduce issues around migration. As such, we have tested a range of lower join ages (75, 80 and 85) against the current ONS value of 90.

Using ONS data from 1972 to 2015 for males, we found that the biggest step change in diagnostic results was between join ages 90 and 85. Two of the test statistics considered, the average annual scaling adjustment and average mortality deviance, are shown in Chart 2.4. We view a smaller average annual scaling adjustment (i.e. closer to 0%) as indicative of better performance, because it means that the KT method has been more consistently in line with the official population totals from year to year. The average mortality deviance is a measure of smoothness of mortality rates at the join age, and we consider a deviance closer to 0% preferable, as this implies a smoother transition of mortality rates when transitioning from the ONS’s census-based population estimates at younger ages to the (constrained) KT method at higher ages. Chart 2.4 shows that reducing the join age from 90 to 85 results in a significant reduction in both the average annual scaling adjustment and the average mortality deviance. Some further improvement in these statistics is seen by reducing the join age further, however, the bulk of the improvement is seen in reducing the join age from 90 to 85.

Chart 2.4: Average annual scaling adjustment (LHS) and average mortality deviance (RHS) at different join ages, KT method with trend and (k,m) = (5,2), ONS male data, 1972-2015

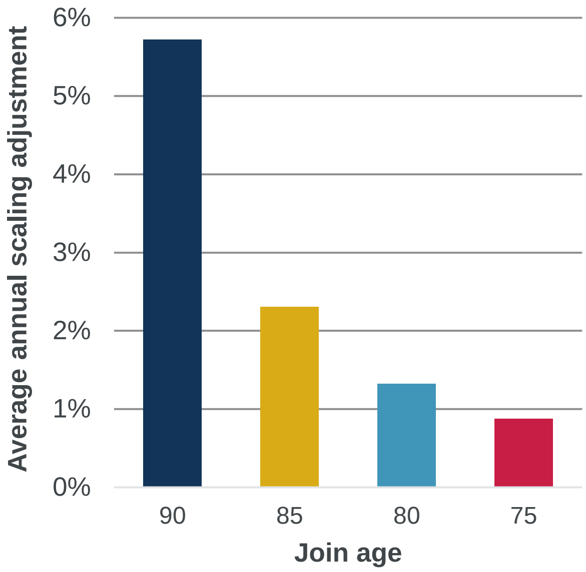

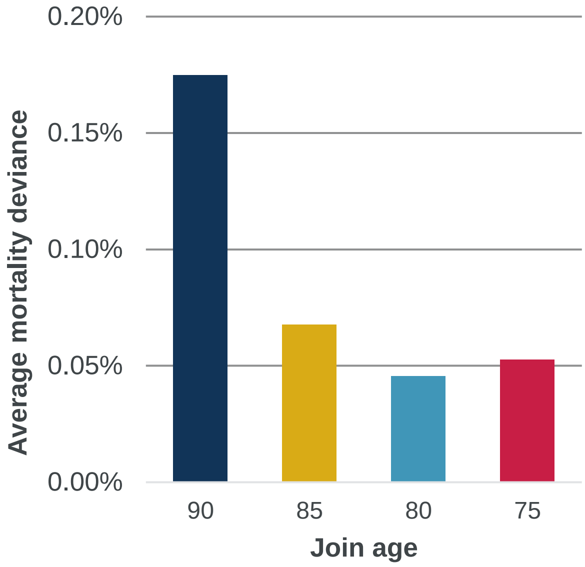

We are conscious that extending the KT method down too far would place heavy reliance on the extrapolation of survivor ratios into ages younger than 90. This means that the method is moving further away from a pure extinct generations approach and into a dependence on mortality improvement projections to estimate the population figures. As join age reduces, there must come a point where our confidence in the stability of the KT method falls below our (increasing) confidence in the official ONS population figures and it is no longer justifiable to override the official figures with estimates.

We note that the bulk of the improvement in diagnostic tests seen in Chart 2.4 comes from reducing join age from 90 to 85, after which only smaller improvements are observed. This may in itself reflect the increasing reliability of the official ONS population figures below age 85 (exhibiting less divergence from the KT estimates).

Bearing all this in mind, we suggest that a reduction in the join age from 90 to 85 is worth considering. It may improve the robustness of the KT method and avoid undue dependence on the current 90+ population constraint, while minimising the greater risks to the reliability of the KT method which a lower join age entails.

Lexis adjustment to deaths data

The KT method requires death counts over a 12-month observation period (typically a calendar year starting 1 January) by age at the start of that period. Contrary to what is required by the KT method, the ONS death registration data is available by age *at death* over each calendar year. The difference between the available and the required data format is illustrated in the so-called Lexis diagram in Chart 2.5.

Chart 2.5: Lexis diagram of deaths data timing

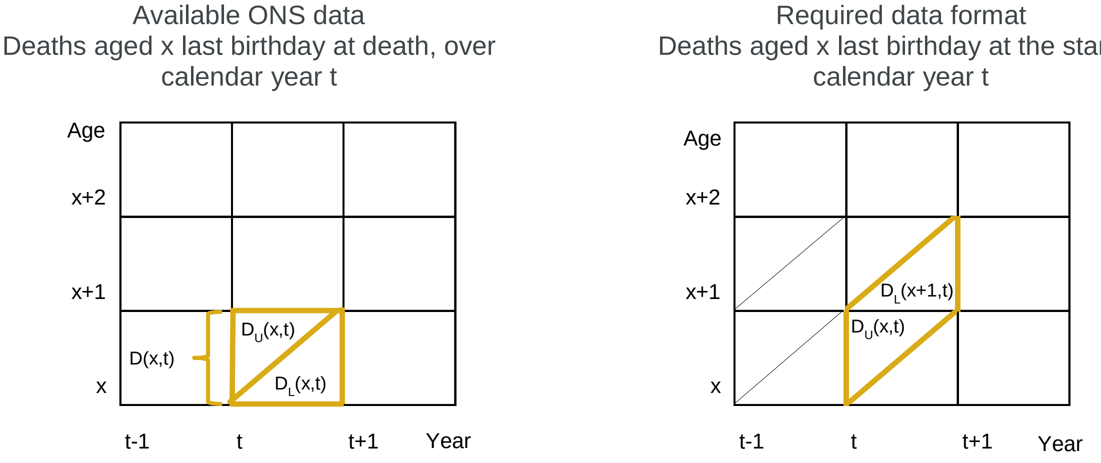

The deaths aged x last birthday at death in calendar year t, **D(x,t)**, can be split into two Lexis triangles (left-hand grid in Chart 2.5). The upper triangle, **D_U(x,t),** represents deaths of individuals aged x at the start of calendar year t, who died before reaching their birthday in the year of death. The lower triangle, **D_L(x,t)**, represents deaths of individuals aged x-1 at the start of the calendar year, who died after passing another birthday in the year of death.

Deaths data as required by the KT method is therefore the parallelogram comprising **D_U(x,t)** and **D_L(x+1,t)** (right-hand grid in Chart 2.5). The challenge is in estimating the proportion of **D(x,t)** that relates to the upper and the lower triangles.

The simplest approach, which is also the ONS method, is assuming a 50/50 split, i.e. the number of deaths aged x at the beginning of calendar year t is approximately 50% of deaths at age x registered in calendar year t and 50% of deaths at age x+1 in calendar year t. However, this assumption quickly becomes questionable at higher ages when the mortality curve gets steeper.

The following refinements have been included in the variant that we have tested:

- Rundown of exposure within a calendar year.
  - Deaths in early months of a calendar year diminish the exposure at later months in the same calendar year.
  - As a result of mortality run-off, exposure for the same cohort is not uniformly distributed over a 12-month period and an allowance has been made to reflect this survivorship profile.
- Unequal cohort sizes.
  - Lives in the upper triangles are aged x at the beginning of the calendar year and older than those in the lower triangles of the same D(x,t).
  - As they belong to different birth cohorts, an allowance has been made to reflect unequal cohort sizes at time t.
- Mortality differential between adjacent ages.
  - As mentioned above, lives in the upper triangles are older than those in the corresponding lower triangles.

  - An allowance has been made so that the monthly mortality rate differs between each cohort, giving a gradual increase between age increments (instead of step changes).

- Uneven patterns of birth distribution.
  - Births by month during a 12-month period are usually assumed to be uniformly distributed.
  - However, this is not a suitable assumption for some birth cohorts, particularly the 1919/1920 cohorts. An allowance has been made so that the monthly exposure within a cohort could vary based on its birth pattern by month.

- Seasonality of deaths.
  - It can be observed from recent ONS monthly death registrations data for E&W that there have been more deaths, on average, in the months January and February than in the months November and December. Indeed, over the period November 2007 to February 2017, comparing January and February deaths with deaths for the previous November and December, January and February deaths were, on average, approximately 9.0% higher.
  - Deaths in January and February are more likely to contribute to the upper Lexis triangle and deaths in November and December to the lower Lexis triangle.
  - We have attempted to allow for this seasonality of death timing in the Lexis adjustments, based on the pattern of smoothed monthly death registration data over calendar years 2000 to 2011 inclusive.

The results of our diagnostic tests, set out in the supplementary Technical Paper, show that the resulting estimates are theoretically more accurate and result in systematic reductions in population estimates at high ages (90+) relative to the current method adopted by the ONS, after incorporating the trend, m = 2, a join age of 85 and the Lexis adjustments. This has a material impact on life expectancy at high ages. The principle of applying Lexis adjustments is a well-established practice already in use by other bodies such as the Human Mortality Database (HMD). Indeed, the HMD describe it as 'one of the most important steps in computing the death rates and life tables'. Therefore, we think there is merit in making allowance for the more sophisticated Lexis adjustments within the KT method.

Exposure adjustments

After incorporating the proposed modifications described so far, there are still idiosyncratic data anomalies apparent in the population estimates (for example, in relation to the 1919/20 birth cohort). The Lexis adjustments partially address these but do not resolve them fully because the Lexis adjustments do not tackle distributional effects in the population and some of the issues relate to data below the join age.

In their paper 'Phantoms never die', Cairns et al considered potential sources of error in the E&W exposure data. As well as proposing a range of diagnostic tests to identify the issues, they also proposed specific adjustments that could be made to resolve them. The CMI Mortality Projections Committee proposed an alternative (pragmatic) approach to smoothing anomalies in the final exposures estimates in CMI Working Paper 91 and have adopted this for CMI_2016. We propose using this simplified methodology to address the remaining data anomalies in the KT method estimate.

The empirical concavity function $C(x,t)$ is one of the Cairns, Blake, Dowd and Kessler diagnostics defined by $C(x,t) = \log m_{x,t} - \frac{1}{2}\left(\log m_{x-1,t} + \log m_{x+1,t}\right)$. This function measures the log-linearity of death rates by age for a given calendar year *t* at age *x* and an improvement would be demonstrated by the concavity function staying closer to zero, without any systematic bias. Chart 2.6 shows that adopting the CMI method of smoothing exposure anomalies results in clear improvements in values of the empirical concavity function of the 1920 cohort, relative to the position without smoothing. Chart 2.6 also shows the "convexity adjustment ratio -1" (the gold dashed line). The convexity adjustment ratio (as described by Cairns et al) is an estimate of how much the mid-year population figure would need to be adjusted to make it more reflective of the exposed-to-risk for each birth cohort. Ideally "convexity adjustment ratio -1" should be close to zero.

Chart 2.6: CBDK – concavity by cohort (1920 cohort) for ONS male data (based on KT method with trend, k = 5, m = 2, join age = 85, Lexis adjustments)

With CMI exposure smoothing from 1975

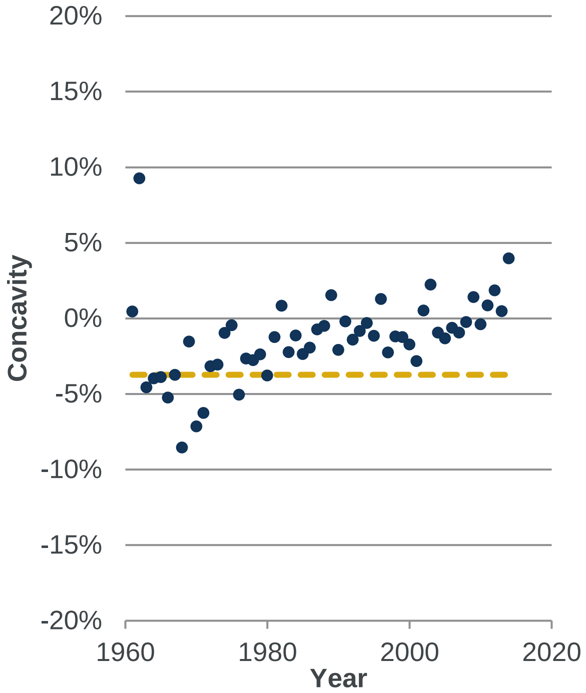

Without CMI exposure smoothing

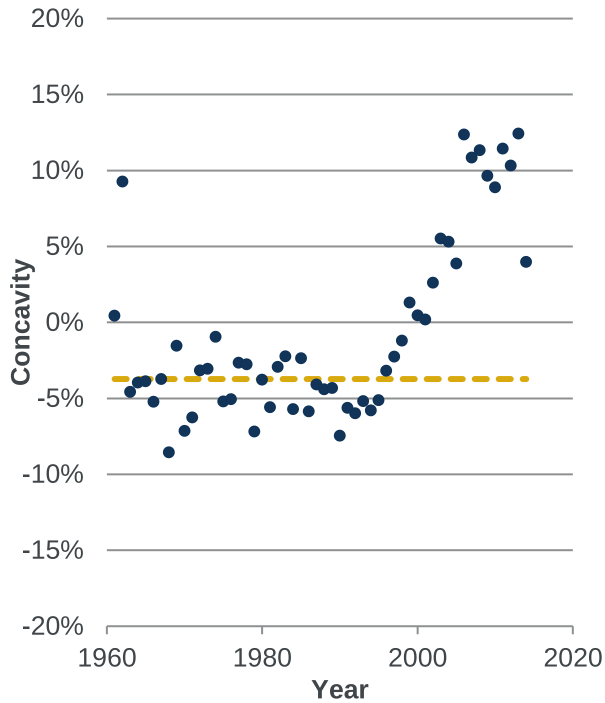

Note that the convexity adjustment ratio shown in Chart 2.6 is a property of the ONS input data only, and does not itself depend on the population estimates produced under alternative variants of the KT method. It simply serves as an informative benchmark to help interpret the concavity in the KT population estimates. For example, it can be seen that the concavity in the unsmoothed population estimates prior to the early 1990s (navy dots on the right-hand chart) is broadly consistent with the convexity adjustment ratio required to correct for the known distribution of births in the 1920 cohort (gold dashed line on the right-hand chart). In other words, the birth distribution anomalies for that cohort largely explain the observed concavity. After the early 1990s there is a further drift in the observed concavity, which may be due to other effects (Cairns et al suggest this may relate to the method used to roll forward the census estimates after 1991). The left-hand chart shows that, after applying the CMI's exposure smoothing, the observed concavity in both regions is reduced.

2.3 Conclusions and next steps

Our analysis suggests that there may be merit in adopting the following modifications to the current ONS methodology for estimating high age population exposures for E&W:

- Refine the projection of survivor ratios in the KT method, by allowing for local mortality trends over time and correspondingly reducing the number of birth-year cohorts (m) in the survivor ration from 5 to 2;
- Extend the high age population method down to a lower age than 90 (for example, age 85) to avoid placing undue reliance on the underlying census-based estimates in this age range;
- Incorporate a more sophisticated approach to adjusting the input deaths data to 1 January age definition (as required by the KT method when using deaths data by age at death) using Lexis triangles;
- Smooth the final exposure estimates as a pragmatic solution to issues arising from uneven birth and death distributions during the year.

Each proposed modification involves varying degrees of complexity, therefore it is possible for readers to consider adopting individual modifications rather than the full package. However, it is worth noting that the modifications of mortality trend and (k,m)=(5,2) should be adopted together. The impact of Lexis adjustment is not very material if the join age is kept at 90, therefore Lexis adjustment is only worth adopting if the join age is lowered. However, the lower join age could be implemented without any Lexis adjustment.

The final (smoothing) element of this proposal is fundamentally targeted at estimating exposure figures *over the year* (as required for mortality analysis), not population estimates *at mid-year* (in line with the ONS's current output). If instead attempting to estimate mid-year population figures, the convexity adjustment embedded in this element would not necessarily be appropriate.

If these changes were adopted, we estimate that the impact on period life expectancy would be as set out in Table 2.3 based on ONS data for E&W over the period 2011 to 2015. Life expectancies are shown in Table 2.3 based on both ungraduated mortality rates (i.e. a simple average of actual death rates in each of the years 2011 to 2015 for each age) and graduated mortality rates (i.e. smooth mortality rates by age that have been fitted using the CMI graduation tool).

Table 2.3: Impact of proposal on period life expectancy, ONS E&W male and female data

Age	Based on average (ungraduated) mortality rates for 2011-2015		Based on graduated mortality rates for 2011-2015	
	Males	Females	Males	Females
85	-0.1%	+0.1%	-0.3%	+0.1%
90	-1.3%	nil	-1.0%	-0.1%
95	-2.1%	-1.1%	-2.1%	-1.0%

The impact could, of course, be different for alternative datasets and time periods and this is a possible area for further research by users interested in this area.

Our analysis of the male data, set out in more detail in the supplementary Technical Paper, suggests that constrained population estimates generated under the KT method with our proposed modifications are broadly higher over ages 85-89 but lower over ages 90-100 than under the base version of the KT method currently used by the ONS. A key driver of this change in shape is the introduction of the local trend allowance, which recognises that survivor ratios have been increasing at a faster rate for younger ages than older ages. The shift in age distribution produces a smaller percentage increase in population over ages 85-89 than the percentage decrease over ages 90-100, because the population size is smaller at the higher ages. This means that the major effect on mortality rates is the increase in mortality above age 90 (and hence decrease in life expectancy at high ages which emerges in Table 2.3 for males).

For females, there is a similar shift in the age distribution, with population decreases at higher ages offset by population increases at lower ages, but the crossover age here is somewhat higher (resulting in a lower life expectancy impact overall).

The supplementary Technical Paper provides further analysis performed on our proposed variants:

- We have considered the impact of updating the population exposures up to 2015 on the latest version of the CMI Mortality Projections Model (the CMI Model). This is calibrated to the ONS E&W data and widely used in the UK pensions and insurance industries to project future national mortality improvements. **We observe small increases in cohort life expectancy as a result of our proposals** (e.g. at age 95, a 0.2% increase for males and a 0.3% increase for females).
- We have repeated a number of our key diagnostics over an alternative period, applying the KT method to data up to 2010 and measuring the impact over 2006-2010. The motivation for this is to check that the performance of our proposals is robust to alternative data periods, and to test the sensitivity of the impact which our proposals have on population exposures (and hence mortality rates). **Our testing has concluded that the changes we are proposing to the KT method in this paper are robust.**

The Working Party will now seek to work with the relevant CMI committees to develop appropriate approaches for their specific investigations.

3. Mortality at very high ages

Working Paper 85 highlighted that there were differing views as to whether the progression of mortality rates into very high ages, i.e. ages 100 and above, follows an exponential (Gompertz) progression or whether there is an element of deceleration in rates. To further this discussion, and to help users of mortality curves to set mortality rates at the highest ages, we have considered three further papers that have studied independent datasets of mortality rates at very high ages, Gavrilov and Gavrilova (2015), Ouellette and Bourbeau (2014) and Rau et al (2016). We set out a high-level summary of each paper below and draw conclusions for what this means for mortality rates at the highest ages.

3.1 Gavrilov and Gavrilova (2015)

In this paper, Gavrilov and Gavrilova (G&G) analyse data from the International Database on Longevity (IDL) which is a project to collate validated deaths data for supercentenarians (those attaining age 110 or more). The data is freely available online at www.supercentenarians.org. The database contains details of 668 super centenarians across 15 countries, with over half of the entries being from the US and 599 in respect of female lives. The deaths broadly occur during the period 1980 to 2007, peaking around 2000. Unfortunately, this database is no longer supported. We have taken the data from IDL as presented and believe it still provides a useful resource of deaths at the oldest ages.

In their paper, G&G took the data and removed deaths of lives born after 1893 due to the risk that these deaths were not from extinct cohorts. The main conclusion from their paper is that the best fitting mortality curve shape for the IDL data is a Gompertz curve (i.e. a linear progression in log μ_x, or alternatively that μ_x is exponential). We have sought to replicate this analysis.

Chart 3.1 plots the hazard rates of all lives excluding those born before 1885 since G&G believe the data for these older cohorts to be less reliable. We have also added high and low gates to show the 95th percentile confidence intervals of these rates which are naturally large, noting that the upper confidence intervals for ages 114 and 115 cap the upper number of deaths at the initial exposure (to ensure that the number of deaths does not exceed the population alive). The results are expressed in log base 10 to be consistent with the approach taken by G&G. The actuarial estimate of hazard rate (μ_x) has been calculated as follows in line with the approach adopted in G&G's paper:

$$\mu_x = \frac{2q_x}{\Delta x(2 - q_x)}$$

Note that we have not assessed the suitability of this formula for estimating μ_x from q_x but have used it purely to ensure a comparable measure with G&G.

Here we can see that the rates from the paper (blue crosses on Chart 3.1) and those we have calculated based on G&G's approach (gold diamonds on Chart 3.1) broadly line up.

Chart 3.1: Replicating hazard rates from Gavrilov and Gavrilova (2015) and alternative hazard rate calculation

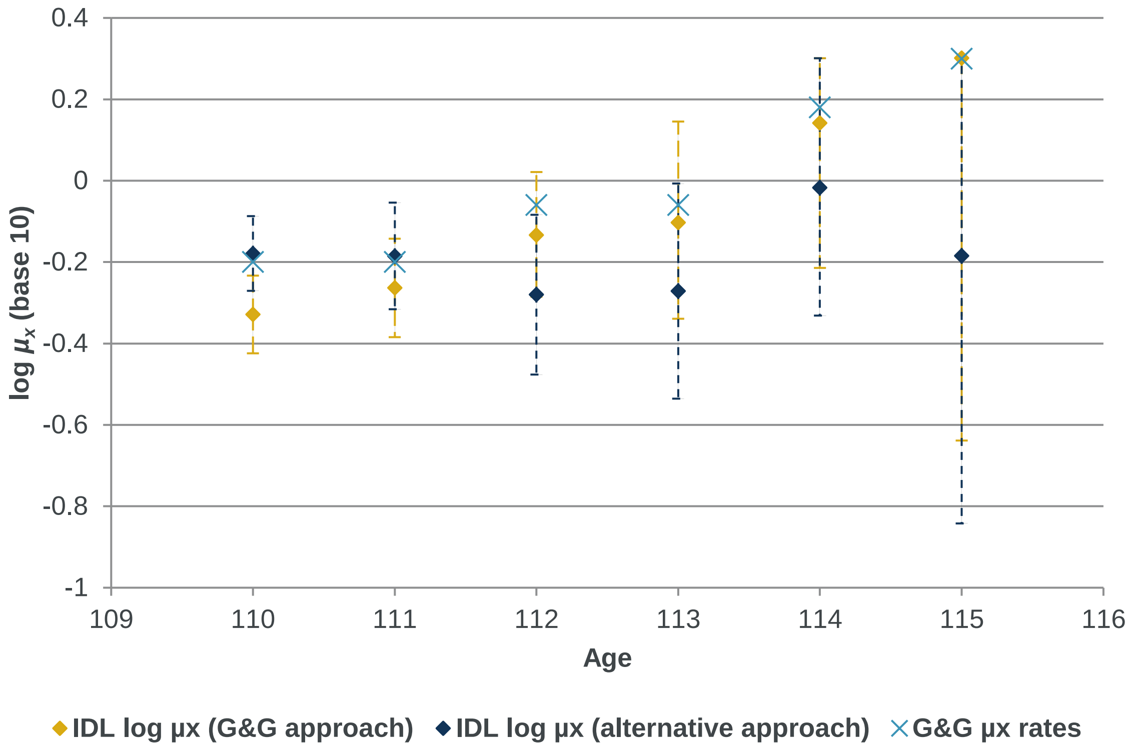

G&G conclude that the best fit curve for the rates as presented above is a straight line and therefore this supports their hypothesis that a Gompertz curve is the best fit for high age mortality, as indicated by the broadly linear plot of the G&G rates in Chart 3.1.

From a review of the IDL data, we have identified the following potential issues and alternative manipulation choices which we are not aware that G&G have allowed for:

- The IDL data commentary sets out start and end dates for the observation periods by country, and we have used these as a more accurate indication of the exposed to risk than simply assuming that each dataset runs up to 2010 (the approach which we understand was taken by G&G). Although we expect the use of these country-specific dates to be more realistic, it is important to note a caveat. Whilst in some cases (e.g. Japan) the period of observation is stated explicitly, in others it is possible that the start and end dates provided simply reflect the timing of the first and last recorded deaths (which means there is a lack of clarity on the actual period of observation for the deaths).
- The commentary accompanying the IDL data acknowledges concerns with the Australian, Belgian and early Italian datasets hence we have excluded these. The concerns regarded inability to verify dates of birth and death for some cases in the Belgian and early Italian data which could mean that ages are mis-stated. The Australian data had suspected age ascertainment bias[1] due to being sourced from media reports, the implication of this being that deaths at more advanced ages are more likely to be reported than deaths at younger ages (which could skew the results of a mortality analysis).
- We have restricted our attention to female deaths only, to try to improve the homogeneity of the data (noting that male deaths represent only a model proportion of the IDL data in any case).
- We have included older pre-1885 birth cohorts[2] to avoid prejudging the data. This means that the earliest birth year we have used is 1858.

[1] See http://www.supercentenarians.org/Home/Data for definition
[2] G&G found that lives from pre-1885 cohorts exhibited a slower increase in mortality than younger cohorts which they suggest could be due to lower quality of birth recording for these cohorts

- Most importantly, we have imposed an extinction age of 118 directly, excluding data for individuals who would not have reached 118 by the country-specific end date. This is in contrast to the approach taken by G&G of simply restricting to birth years before 1893, which is less certain to ensure cohorts are extinct by the end date of observation (for some countries, lying many years prior to 2010). It is important to ensure that all cohorts under analysis are extinct, in order to avoid missing exposure for those lives who died after the country-specific end date. The analysis is very sensitive to this difference in approach, especially at the highest ages 114 and 115.

Chart 3.1 shows the impact of making the above changes in the data manipulation (blue diamonds labelled alternative approach) and is otherwise comparable to the G&G approach. One can see that the conclusion of a straight line progression is less obvious.

In light of the identified wide confidence intervals, we think that it is only possible to make inference about the shape of the curve at very high ages by putting these rates into context with those from younger ages. The analysis set out above only considers six age points, in practice the shape at younger ages will also help to provide further insight.

3.2 Ouellette and Bourbeau (2014)

Ouellette and Bourbeau study the death rates of Canadian centenarians using Quebec's church parish registers. The authors were able to match birth dates on death certificates with baptismal certificates for the majority of cases studied which gave a high degree of certainty around the data quality. The study includes 1,712 lives.

Their paper seeks to verify / challenge the claim in Gavrilov and Gavrilova's 2011 paper that there are a number of factors that explain observed mortality deceleration, and that when controlled for, deceleration is not observed in the data under consideration. In summary, these were:

i. Aggregating birth cohorts leads to heterogeneous groups and therefore an apparent slow down in the mortality curve is observed as the more frail lives die at younger ages. This effect is discussed in Section 4.1 of Working Paper 85.

ii. The standard assumptions regarding hazard rate estimates, in particular the assumption of a uniform distribution of deaths within each one-year age interval, may be invalid when the risk of death is extremely high at old ages and thus lead to a slowdown in the hazard rate increase at very old ages.

iii. Inaccurate age reporting leads to a downward bias in hazard rate, as discussed in Section 6.1.5 of Working Paper 85.

To address (i), the authors show that, for their dataset, the fitted hazard rate varies very little if the data is split by birth cohort (Figure 4 of their paper). The authors also make the point that their data is relatively homogeneous in that it is from lives born in Quebec only.

For (ii), the authors imply that Gavrilov and Gavrilova have erroneously multiplied probabilities of death for periods less than a year to derive hazard rate estimates that allow for the non-uniformity of hazard rates at advanced ages and provide an example of the results one gets from doing this (in Figure 3 of their paper). However, in a response to this claim[3], Gavrilov and Gavrilova are clear that they were not multiplying probabilities of death but working with hazard rate estimates. It is unclear to us, without seeing the actual calculations, whether the Ouellette observation is fair or not.

For (iii), Ouellette and Bourbeau demonstrate that they have performed vigorous validation on the data they have used to gain comfort that the ages at death are accurate.

The result of this analysis is that there is observed deceleration in the death rates at older ages and we have plotted the crude annualised m_x rates from their analysis in Chart 3.2 (see Section 3.4).

We note that in their response to this paper, Gavrilov and Gavrilova comment on the low number of deaths in the Ouellette study (for example seven deaths after age 109) which we believe is a fair criticism and to that end we have included confidence intervals on the death rates from the paper. The authors acknowledge that a dataset

[3] http://longevity-science.blogspot.co.uk/2016/01/rebuttal-of-criticism-of-our-longevity.html

with deaths from lower ages would be preferable to support estimation of the shape of mortality rates at the higher ages.

Our view is that this paper is a helpful independent view on death rates at advanced ages on a relatively clean dataset and so can be used as a guide to mortality level. However, inferring the mortality slope from this dataset alone is likely to be prone to mis-estimation errors due to the lack of data and reliance on one territory.

3.3 Rau et al (2016)

This paper was presented to the 2017 "Living to 100" conference and seeks to verify the results of Gampe (2010) by looking at the age and level of mortality plateau across a number of developed countries using data from the Human Mortality Database. As part of this study, the authors fitted Gompertz, Gompertz-Makeham, Gamma-Gompertz and Gamma-Gompertz-Makeham models to male and female data from seven countries and across different time periods. The Gamma models refer to frailty models as referenced in Working Paper 85 where the Gompertz curve level is multiplied by a random variable with a gamma distribution to represent the heterogeneity in the population (see Missov and Vaupel 2015).

For the Gamma-Gompertz model, the hazard rate is given by the following formula, where the authors are expressing the Gompertz form in terms of the modal age at death as set out below. Under the Makeham models a constant is added to the hazard rate.

Gompertz: $\mu(x) = be^{b(x-M)}$

Gamma-Gompertz model: $\mu(x) = \frac{be^{b(x-M)}}{1+\gamma e^{-bM}(e^{bx}-1)}$

The Gamma-Gompertz model will exhibit deceleration as age increases and will tend to a plateau which is given by:

$$\lim_{x\to\infty} \mu(x) = \frac{b}{\gamma}$$

The authors found that the best-fitting model (by the AIC criterion) across virtually all time periods and countries was the Gamma-Gompertz or Gamma-Gompertz-Makeham; that is the models in which the mortality curve decelerates to a plateau. The best fitting value of γ, the variance of the Gamma distribution was around 20% for females. This implies a standard deviation of 45% in mortality level which seems to be plausible given this relates to differences between individuals rather than pension schemes or annuity blocks. Note that Gompertz models are a special case of Gamma-Gompertz models where γ is equal to 0.

The authors found that the level of the plateau for μ_x was around 0.8 for females with inter quartile range of [0.7 – 0.9]. For males the result was around 1.2 although the male results showed significant variation (interquartile range of [0.8 – 1.6]). They also showed that they were able to narrow the range of plateau results for females after restricting to six[4] large, low mortality countries with high quality data, although this approach did not reduce the variability for males.

Rau et al conclude that there is good evidence supporting a plateau of around 0.8 for females. This is very much in line with the results of Gampe (2010) which proposed a plateau of 0.7 and was based on predominantly female data. They also note that if there is a plateau for males it is higher but that their estimates were not convincing. A key takeaway is that the variation in the plateau estimates for men is driven by the γ parameter, suggesting that the level of heterogeneity in male data between countries is potentially wider than females. However, as the authors comment, we need more male deaths at advanced ages to support any concrete conclusions. As far as we are aware there are no other papers that have tried to calculate the level of the plateau for males.

We think this paper is a very helpful contribution to the debate on the shape of the mortality curve at high ages since it analyses a very wide range of countries through time as opposed to focussing on a particular country and specific birth cohorts. This also has the advantage that particular data recording or quality issues specific to a single country are not skewing the results.

[4] Belgium, France, Germany, Germany-West, Italy and Japan – all have a population size of greater than 10 million for the years 2005-2010.

3.4 Implications for mortality rates at very high ages

Chart 3.2 summarises our analysis of the contributions of the three papers. We have plotted the raw m_x rates from female E&W ONS data to age 113 (incorporating the adapted KT methodology proposed in Section 2) labelled "E&W", the raw m_x rates from the Ouellette and Bourbeau paper, and our revised m_x rates from analysis of the IDL database, with 95% confidence intervals shown for the Ouellette and Bourbeau and IDL m_x rates. For reference, we have also plotted a Gompertz curve (G(2)) and Gamma-Gompertz curve fitted to the adjusted E&W data. The Gamma-Gompertz curve has been fitted as per the specification in the Rau et al paper and the implied mortality plateau from the fit for μ_x is 1.06, tending towards this level at around age 150. The value for m_x around age 120 is very close to 1.

Chart 3.2: Combined E&W 2005-2010 and IDL mx rates with graduated curves

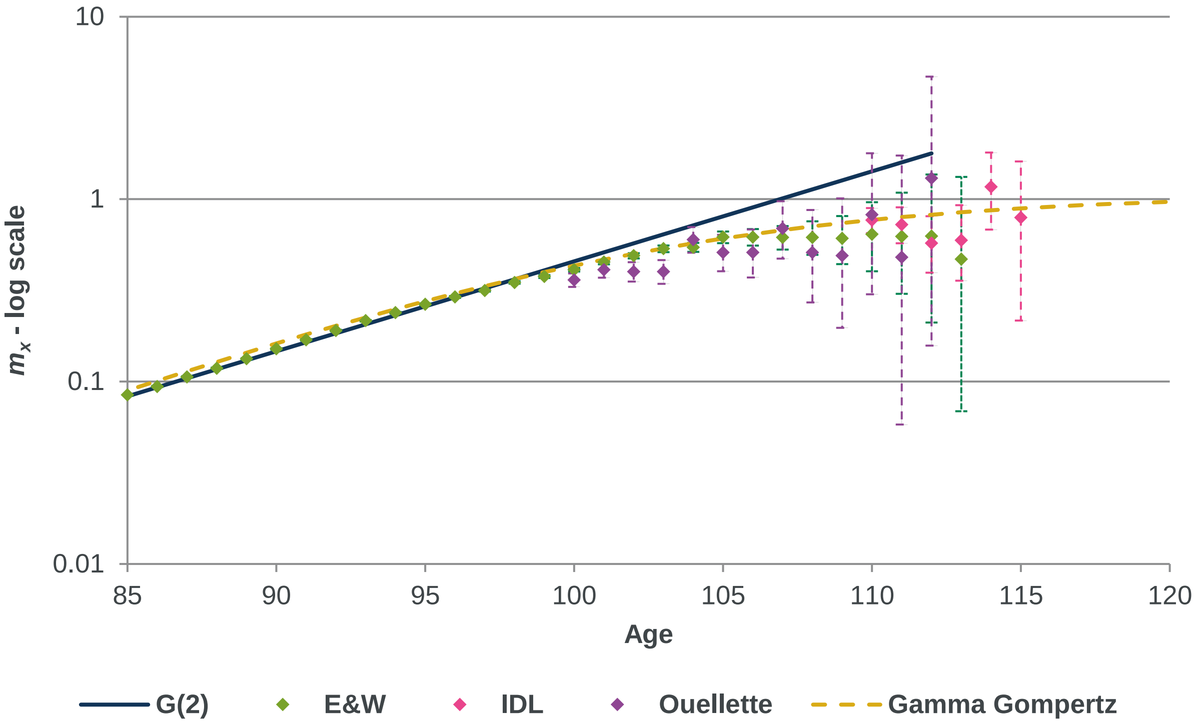

Our conclusions from this graph are that it is hard to justify a log-linear fit to the E&W data in light of the data sources we have collected and the confidence intervals around them. Furthermore, the Ouellette and IDL m_x rates seem to be broadly consistent with the fitted Gamma-Gompertz model when applied to the E&W data. The Gamma-Gompertz model also seems to support a value of μ_{120} of around 1 which is the value currently used by the SAPS and Annuities committees. As set out in Section 4.2 of this paper, the impact of allowing for deceleration is a 2.3% increase in life expectancy at age 95 instead of fitting a G(2) curve to the female data.

With respect to the Gavrilov and Gavrilova proposition of a log-linear progression of the mortality curve, one helpful observation made at the 2017 Living to 100 symposium by Jean Marie Robine was that the G&G analysis (in particular the G&G 2011 paper) looks at the mortality curve on a cohort basis which necessarily will produce shallower gradients compared with a period approach due to the material impact of mortality improvements over the last century, hence any deceleration at the advanced ages will be less pronounced.

Based on the papers we have reviewed and analysis we have performed, we believe that a mortality curve which makes allowance for deceleration at advanced ages is supportable for period mortality at population level, in particular E&W. We also believe that a value of μ_{120} of 1 is justifiable and a reasonable assumption.

4. Framework for setting high age mortality rates

In this section, we set out our proposed framework for setting mortality rates at high ages. This has been written in the context of determining CMI mortality tables but we hope it will be helpful in other contexts. We start by considering desirable features, and then set out our approach.

4.1 Proposed approach to setting mortality rates at high ages

Desirable features

In this section, we set out the desirable features for setting mortality rates at high ages (90+) that we have considered:

- Compatibility with data. Where there is mortality data at high ages of appropriate quality, quantity and relevance then this should be considered to inform the rates. Where the quantity or quality of this dataset is being questioned then alternative datasets and/or expert judgement may be required to determine an appropriate modelling approach, accepting that the direct relevance may be reduced.
- Plausibility. For high age mortality, this is intended to cover both objective analysis (where there is data to inform the level and shape of mortality), as well as subjective criteria related to biological, medical and environmental forces. For the latter case, a modeller of high age mortality may have to consider what plausible combination of factors would have to occur to produce the assumed mortality. If the modeller cannot determine a plausible scenario, then the proposed approach should be ruled out.
- Allowance for features for particular groups of lives. We recognise that at younger ages, different groups of lives exhibit differing mortality experience across the same age range. We have considered whether the mortality for different cohorts converges at very high ages and if so, how this could be allowed for when producing mortality tables.
- Allowance for trends in mortality. Where historical data has been used to calibrate a graduation, allowance for mortality improvements between the analysis date and the table effective date should be allowed for.
- Smooth progression. We expect a smooth transition between the portfolio mortality rates and the rates determined for very high ages

Deconstruction of high age ranges

In developing a framework to determine the mortality curve at high ages we have split the age range into three categories.

- Age range A – Over this age range, graduated mortality rates are set principally by reference to the experience of the portfolio in question. As such the question of convergence over this age range is driven by the experience data.
- Age range B – Over this age range the portfolio data is not sufficient (in volume and / or quality) to reliably graduate mortality rates, and so wider data (e.g. relevant population data) needs to be used to inform the graduation / extension. A view must therefore be formed on the degree and way in which there is convergence between the rates being graduated and the population data in this age range.
- Age range C – Over this age range there is no portfolio data and limited population data available. Here the user should consider applying expert judgement informed by any available data as well as plausible views as to the shape and level of mortality. This approach will in turn imply an assumed ultimate age and mortality rate at the maximum age of the mortality table.

Chart 4.1 summarises this division of age ranges. The age ranges will vary depending on the dataset and purpose of the graduation. The availability, reliability and credibility of data will inform the maximum ages for each of the three age ranges.

Chart 4.1: Defining the age ranges

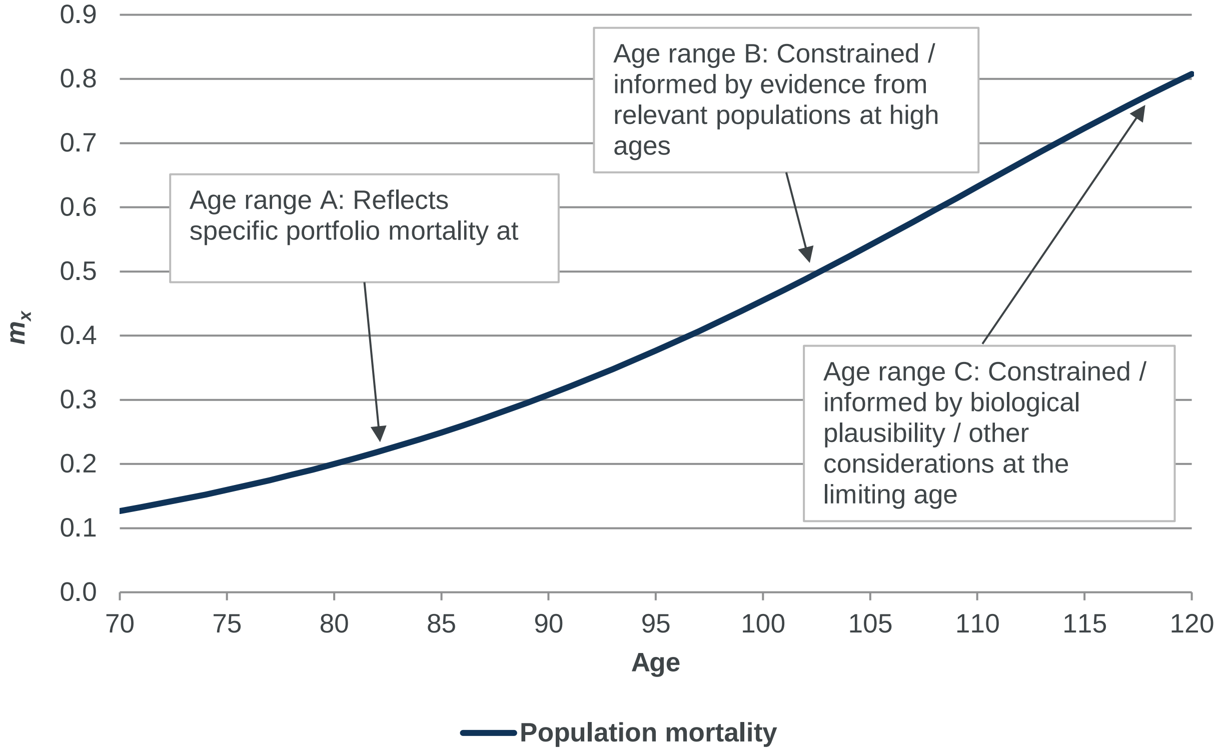

Proposed framework

We break the framework down into individual steps as follows, providing illustrative example charts for each step. This section is intended to be high-level and principles-based. We intend to work with CMI mortality-related committees to support the implementation of this framework for future mortality tables.

Step 1: Graduate Portfolio data

We expect to have available portfolio data informing the graduation at younger ages. Where there is available and reliable data, we expect the graduated rates to be set with reference to the portfolio data, with judgement applied as to the age range where this is applicable. There will be an age reached where the data is deemed to be less reliable.

This maximum age can be determined by an analysis of the data and statistical tests (which are not the focus of this paper), and following graduation we may have a position as illustrated in Chart 4.2.

Chart 4.2: Graduating portfolio data

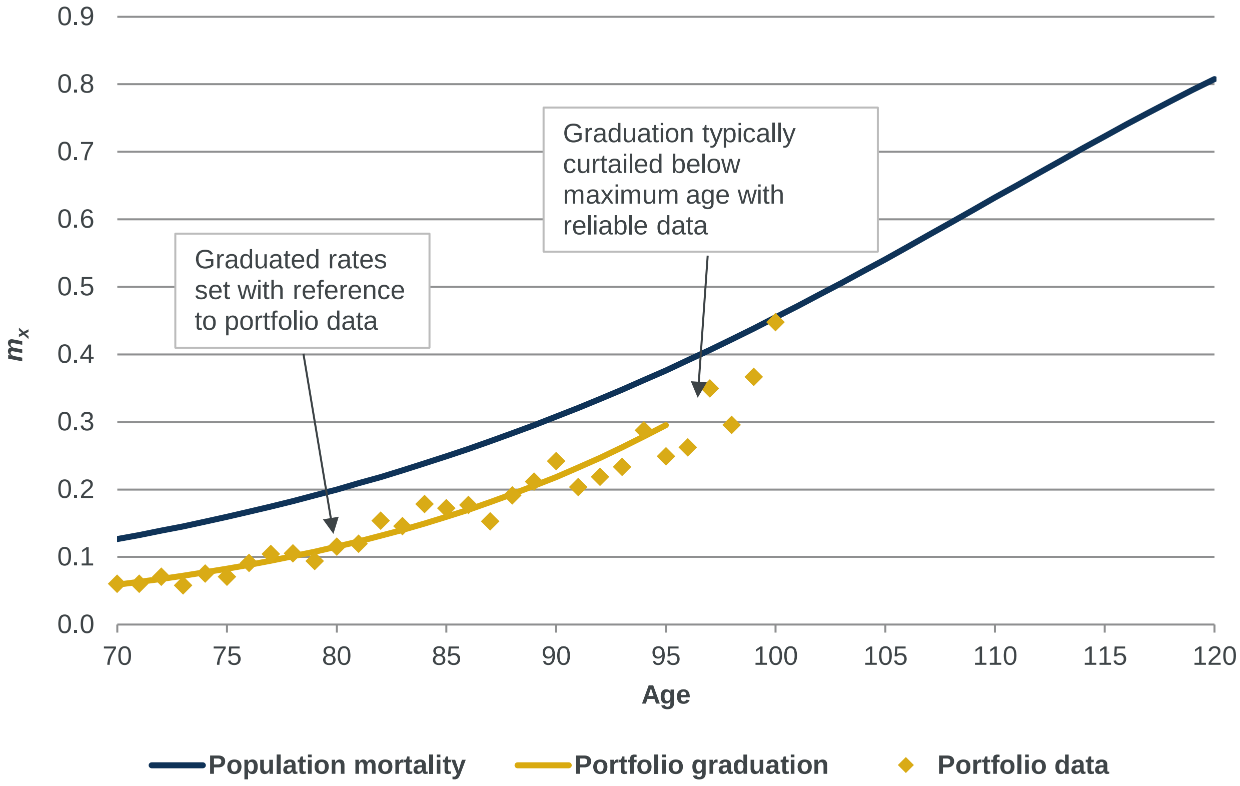

Step 2: Analyse convergence

We then consider the trajectory of portfolio data relative to population data. In particular, statistical measures of the trajectory of portfolio to population data can be considered where both sets of data have sufficient volumes. The metrics for analysing convergence are considered further in Section 4.3. It may also be possible to consider relevant wider data, such as industry data, where portfolios are immature, as a stepping stone between portfolio and population data, where such data is available and of sufficient quality.

Where there is evidence of convergence and this is believed to be a genuine feature, this analysis can be used to inform an extension as we move from using portfolio data to population data to set the mortality rates at higher ages. The questions we may typically ask in this analysis are shown in Chart 4.3.

Chart 4.3: Assessing convergence to population data

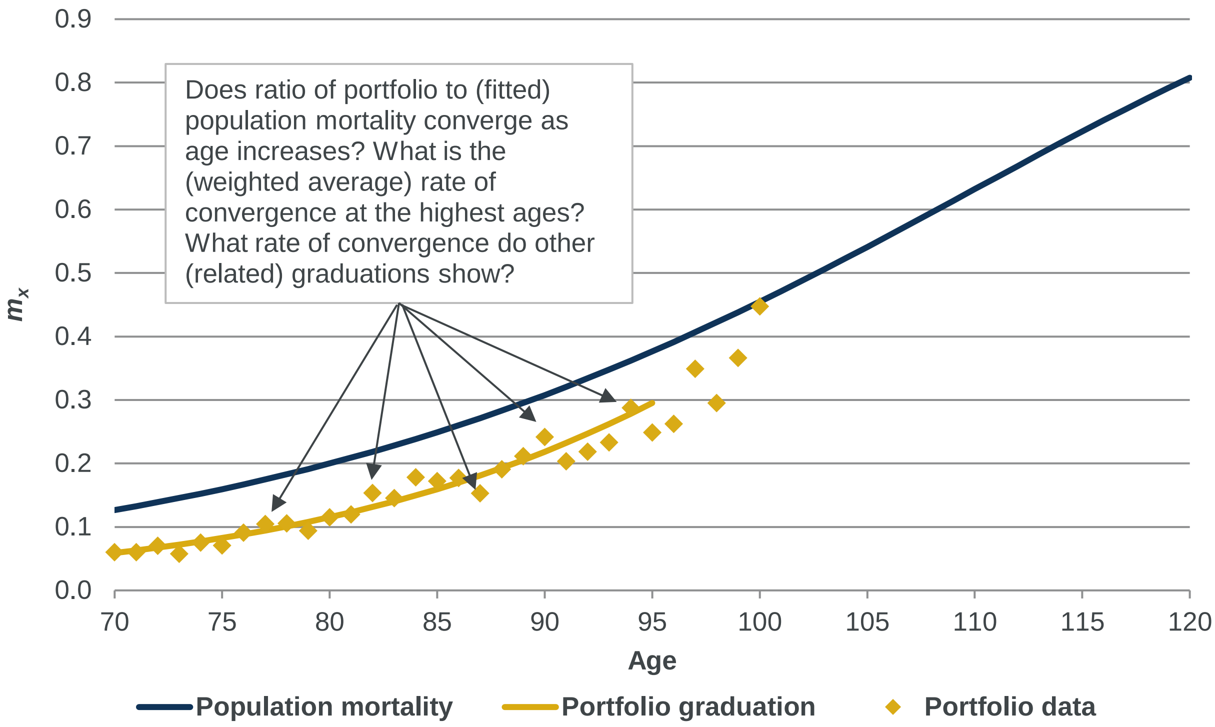

Chart 4.3 shows a stylised picture, with smooth convergence in underlying mortality rates (when considered in relative terms) albeit with stochastic variation around the observed rates. In practice, a more complex pattern may be seen, for example divergence at younger ages, followed by convergence at higher ages. In analysing patterns to inform the extension we would suggest focussing on the age range above which there is a reasonably consistent pattern of convergence (or indeed non-convergence if that were the case).

Where there is no evidence of convergence, then the user should consider whether convergence should be applied or whether a difference between portfolio and population data should be maintained.

Step 3: Extend graduation

We can now extend the graduation, converging from portfolio to population data.

Volumes of data reduce rapidly as the age range progresses from B to C. Consideration needs to be given as to how to close-off the mortality table. The graduation from population data should help inform the limiting age and rate. There may also be information of mortality at very high ages to inform the level of rates beyond the level at which population mortality data volumes are limited or not available. The analysis presented in Section 3 will help inform the level and shape of mortality at the oldest ages.

Chart 4.4 illustrates this extension of the graduation. Under the proposed framework this is achieved by extending the portfolio graduation allowing for convergence towards the population graduation, at a rate set having regard to the convergence observed in step 2 above.

From a practical perspective, there can be merit in assuming the same rate of convergence to the population graduations for related tables (provided this is not inconsistent with the analysis from step 2). Doing so can avoid introducing crossings in the extensions of related tables.

Chart 4.4: Extending graduation to ultimate age

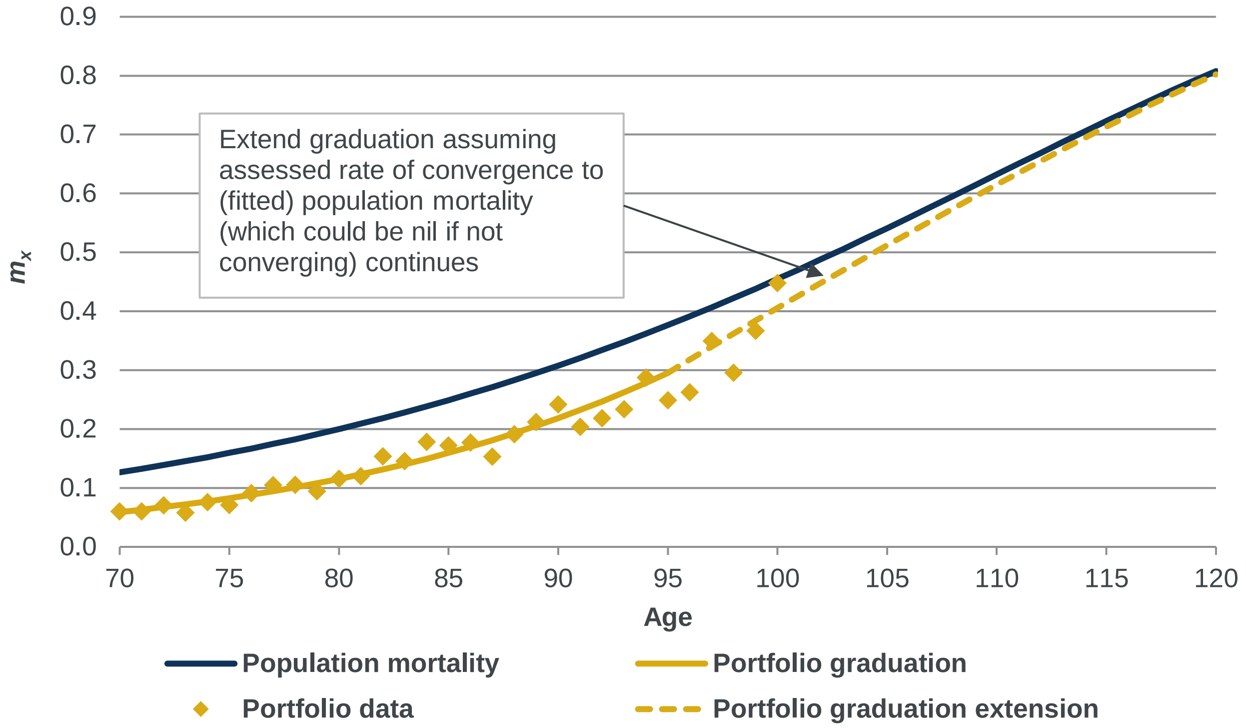

We can extend this framework to cover the graduations of different sub-portfolios (child) relative to a primary (parent) graduation. This is conceptually the same as the above approach.

For example, the SAPS tables have graduations varying by high and low pension bands, alongside the main pensioner graduations. Any graduation of a dataset that is a subset of the main graduation (i.e. the parent table) could be considered a child graduation and should be allowed to converge into the parent table at high ages, if convergence is observed to occur. The key objective here is to ensure consistency between the parent and child graduation. Where there is an assumed convergence profile, the level and shape of rates for the parent and child graduations should also be consistent, avoiding any unsubstantiated anomalies such as cross-over of mortality rates.

Chart 4.5 sets out the extended approach. A similar approach is initially adopted where mortality data is available for the child portfolio. Graduated rates are set with reference to the child portfolio, with particular focus on where the graduation should be curtailed as the quality and credibility of data drops off. The point at which data volumes drop off for the child portfolio will be expected to be at a lower point than the adult portfolio given the lower data volumes, although it may be similar depending on the relative age mixes of the two portfolios.

Chart 4.5: Application of methodology to sub-portfolio data – graduating sub-portfolio data

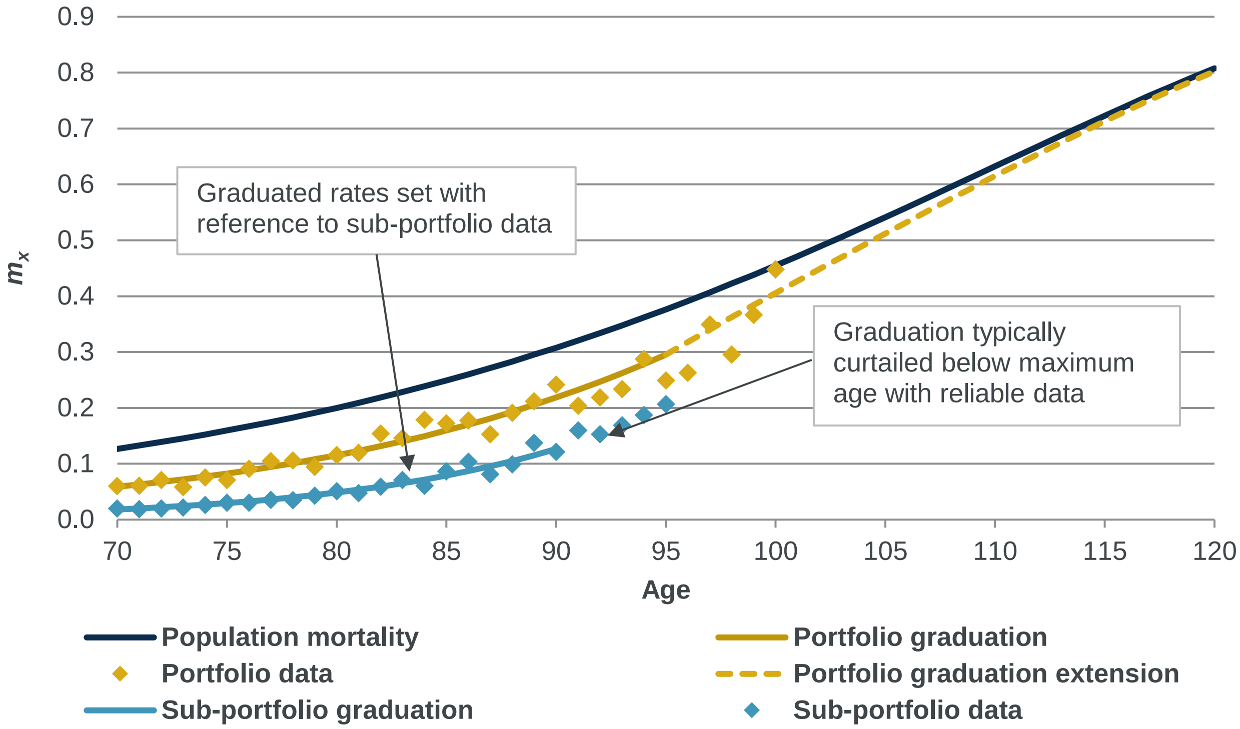

Whereas convergence towards population mortality has been assumed for the parent portfolio, the rate of convergence for the child portfolio may be analysed relative to the adult graduation, as shown in Chart 4.6.

Chart 4.6: Application of methodology to sub-portfolio data – assessing convergence

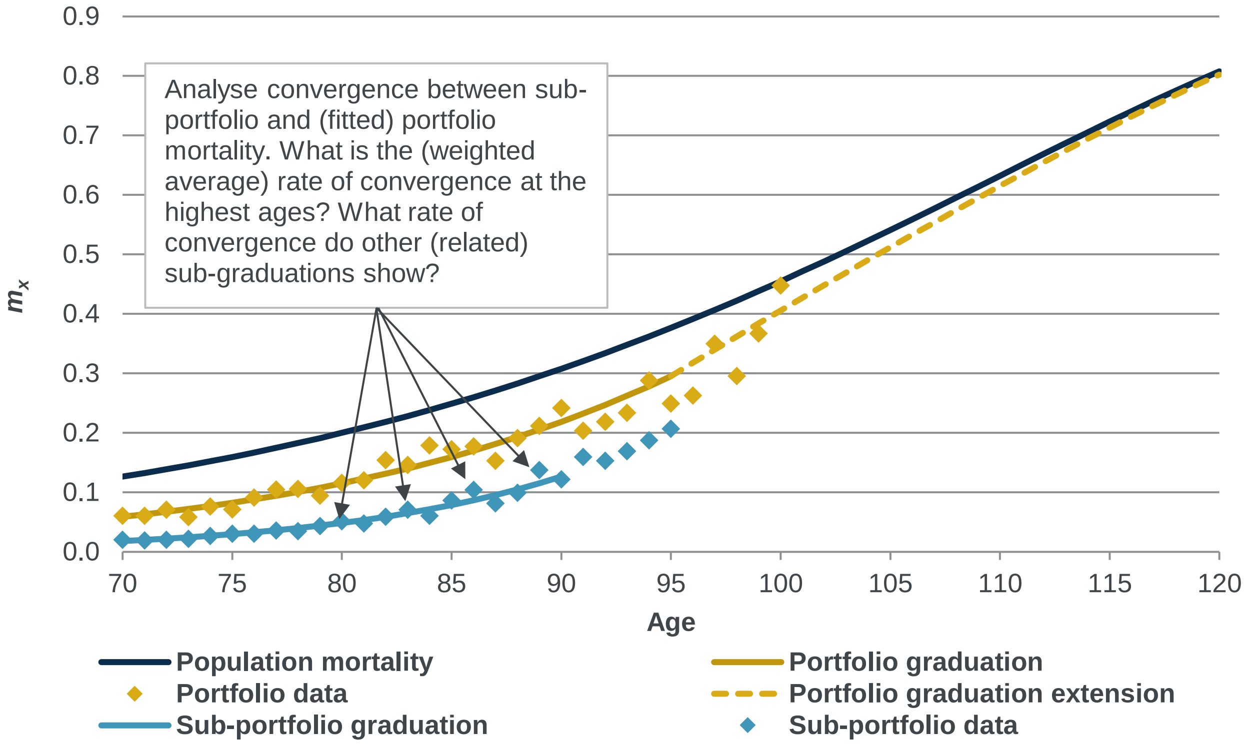

Once judgement has been applied around the rate of convergence for the child portfolio, then the graduation can be extended to older ages relative to the adult portfolio graduation as shown in Chart 4.7. In particular, there may be a desire to have the same mortality rates for both portfolios at very high ages, i.e. at ages 100 and above, or perhaps from a slightly younger age if supported by the data. The child rates will tend to the parent rates, with the same ultimate mortality rate achieved at age 120. Users should consider whether rates should be equal from a younger age than age 120.

Chart 4.7: Application of methodology to sub-portfolio data – extending graduation

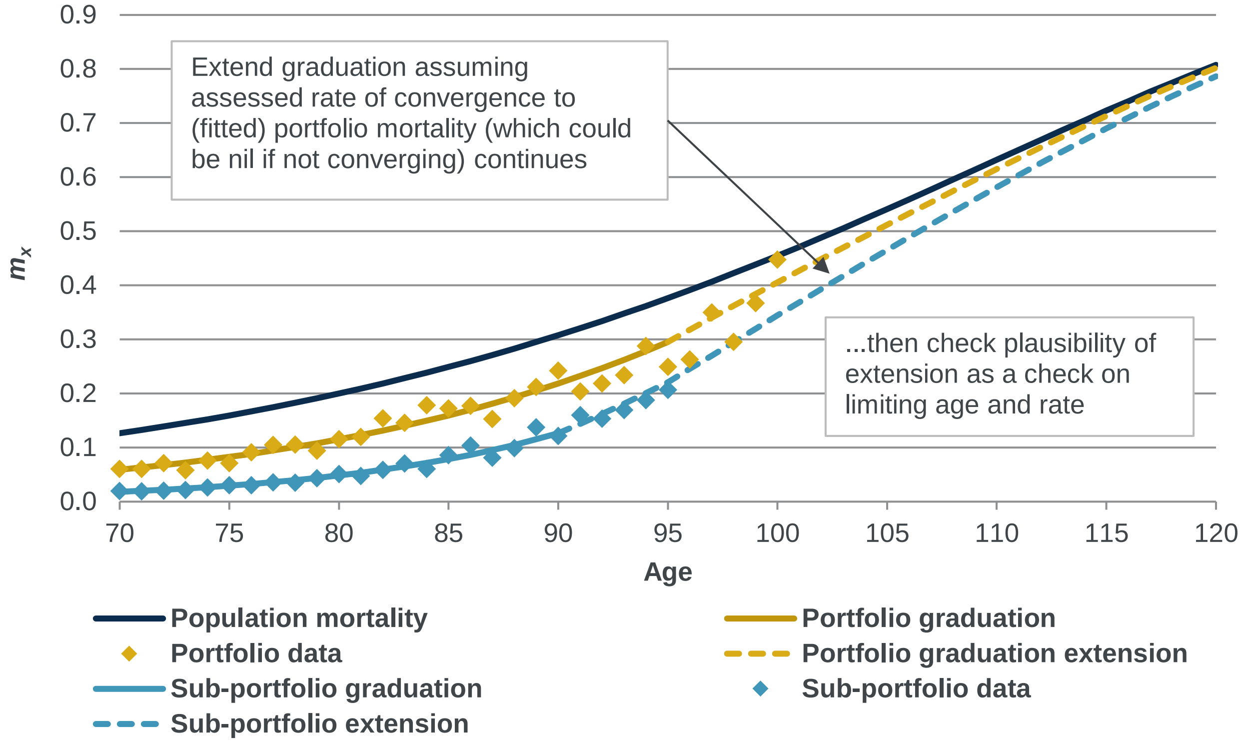

Assessment of proposed framework

In Table 4.1, we assess the proposed framework against the desirable features set out in Section 4.1.

Table 4.1: Assessment of framework against desirable features

Desirable Feature	How addressed
Biological plausibility	Applied to determine mortality at very high ages (age range C), where there is a lack of quantity and/or quality of data. Mortality at very high ages was considered further in Section 3.
Compatibility with data	Applied for age ranges A and B, where graduations are set with consideration of available portfolio data and later population data when portfolio data volumes and/or quality of data is reduced. Examples of the application of this methodology are set out in the case studies in Section 4.3.
Allowance for features for particular groups of lives	We address this feature through considering whether mortality experience converges for different sub-portfolios. Specific examples are considered in the case studies in Section 4.3.
Allowance for trends in mortality	In our testing, we have used data from consistent periods to allow for the presence of trends in mortality.
Smooth progression	We consider the smooth progression of the graduated extended mortality rates in the testing.

4.2 Determining a suitable level of population mortality (Age Range B)

In this section, we consider appropriate mortality rates to use in age range B. A population dataset is by definition likely to have sufficient data in age range B to provide credible mortality rates against which to benchmark mortality rates from the portfolio dataset. The population dataset should be chosen to cover a reasonably similar time period to the portfolio dataset but, where there are timing differences, projections should be applied to the population table to bring the timing of the two tables in line.

In our case studies, in Section 4.3, we have leveraged ONS data of deaths by single year of age and calendar year as well as population estimates as estimated by the proposed methodology set out in Section 2 of this paper (referred to as New KT). Producing our own estimate of population mortality has allowed us to use data for years broadly consistent with those underlying the SAPS (2004 – 2011) and annuities (2007 – 2010) most recent graduations. We have graduated population data over the period 2005 – 2010 to be broadly consistent with both tables and, given mortality improvements at high ages have been relatively low, we believe this is appropriate for illustrative purposes. For actual graduations, a consistent period would obviously be preferred. It should further be noted that our population table is necessarily lives-weighted while the majority of portfolio graduations may be amounts-weighted. We consider these differences to be minor at the ages being considered, due to observed convergence of mortality rates by pension amount at higher ages.

Based on the analysis in Section 3, we have used a value of μ_{120} of 1 and have applied graduations which allow for deceleration away from the Gompertz (G(2)) form, in line with our findings from Section 3.

We have adopted the following principles when constructing the graduation and extension of tables, motivated by the desirable features we have set out in Section 4.1:

- Fit a polynomial to the log of m_x and choose the order of the polynomial by considering a number of criteria including minimising AIC / AICc / BIC as defined in in the CMI Graduation Software User Guide;

- Use as much of the underlying data as possible, i.e. only reject data points if they are implausible;
- Use a high age extension to set the shape of the curve at very high ages and to maintain a smooth progression of mortality rates.
- Set run-in age (i.e. the age at which the population table begins to converge towards the ultimate mortality rate at the final age included in the table) as high as possible and choose a curvature parameter to maintain smooth progression of mortality rates to μ_{120} =1.
- Consider sensitivities in life expectancy for various graduation choices in order to assess materiality.

If using an existing population table rather than constructing graduation for convergence analysis, care should be taken in understanding the ultimate mortality rate, i.e. μ_{120}, and, if the ultimate rate is not in line with the user's beliefs about mortality rates in age range C, then it might be necessary to consider using an alternative population table or blending approach.

Base graduation

We set out in Charts 4.8 and 4.9 graduations for E&W females and males respectively for different order polynomials as well as showing the crude mortality rates. We have used all available high age data and so have considered data over ages 60 – 113 for females and 60-110 for males.

Although the G(7) has the best information criteria performance for both males and females, we have selected a G(5) graduation on the basis that a simpler form is more desirable and that the G(5) graduation is less responsive to the slightly more volatile m_x rates at very high ages when we have higher uncertainty. We also found that apart from the G(2) and G(3) graduations, the life expectancies from the graduations are all very similar up to very high ages. In light of the very low observed value of m_{110} for males, we have chosen to use an age range of 60-109 for the base graduation for males (albeit this choice makes very little difference to the overall result).

Chart 4.8: E&W Female 2005-2010 graduated mortality rates

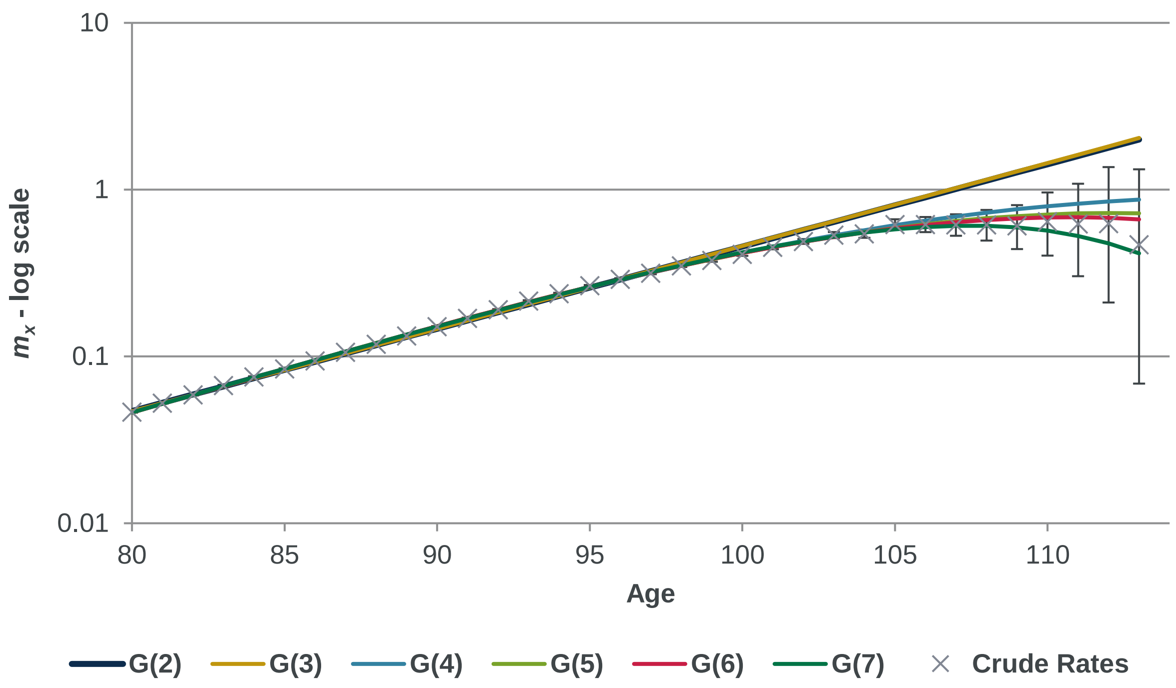

Chart 4.9: E&W Male 2005-2010 graduated mortality rates

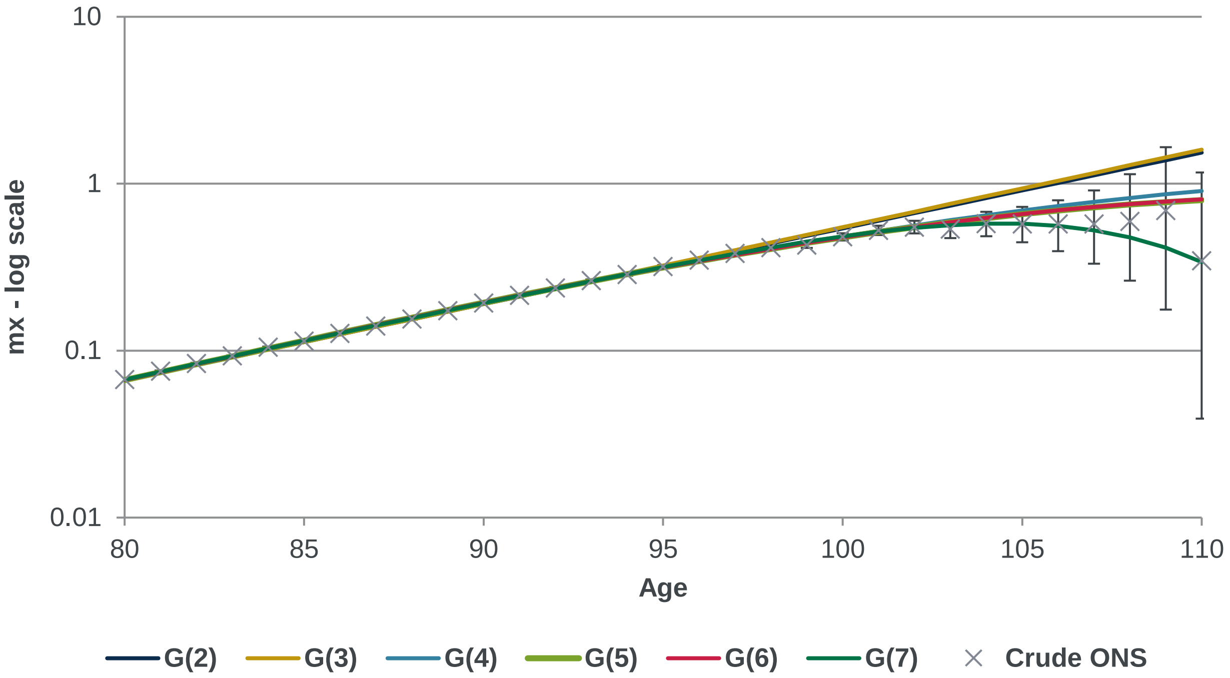

High age extension

There are four parameters required for the high age extension methodology as used in Working Paper 78: the run-in age, curvature parameter, highest age and value of μ_x at the highest age. Given we have set the highest age at 120 and want to use a value of $\mu_{120} = 1$, we only need values for the run-in age and curvature parameter. We have used a run-in age of 105, motivated by our objective to make this as high as possible but constrained by our desire for a smooth progression of mortality rates.

To illustrate this, we have set out three permutations of the female graduation in Chart 4.10. The dark blue line graduates the data up to age 109 then adopts an extension run-in age of 105 (i.e. looking at the impact of excluding the highest four graduated data points). The gold dashed line looks at graduating to age 113 with a run-in age of 105 and the lighter blue line graduates to age 113 with a run-in age of 113. Although the life expectancies of these permutations at age 95 are all identical to two decimal places, clearly the lighter blue line fails the test of smooth progression of mortality rates. We have also looked at the sensitivity of using a run-in age of 100 and see little impact on the results in terms of shape of curve or life expectancies. We have therefore chosen a run-in age of 105 for both genders and graduated to age 113 for females and 109 for males. The approaches adopted for SAPS and Annuities data is discussed further in the case studies that follow in Section 4.3.

We have chosen a value for the curvature parameter of 1.25 which supports the smooth progression of the mortality curves and we have set this by eye (however, given the choice of run-in age, the value of this parameter is not very material).

Chart 4.10: E&W Female 2005-2010 graduated mx with high age extension

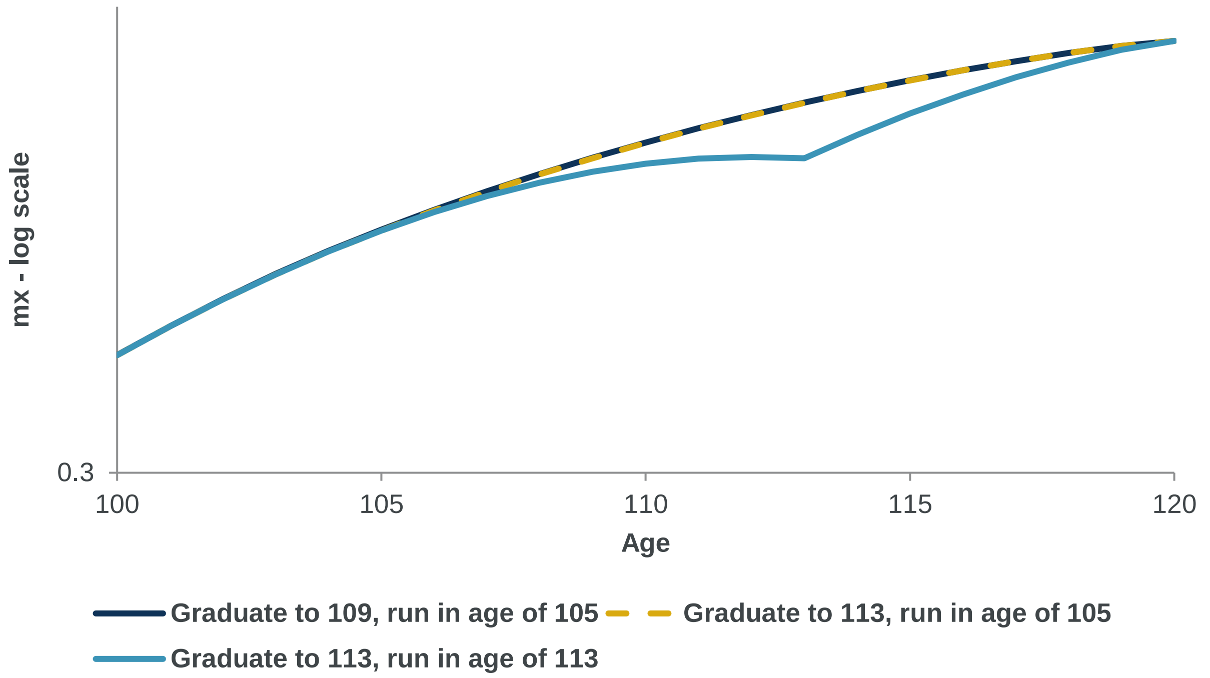

Graduation sensitivities

To test the materiality of the graduation choices made above, we have considered the following sensitivities to the graduation approach:

- Use a G(2) graduation instead of G(5);
- End the graduation age at earlier ages (100,105,109 for women and 100,105 for men);
- Change the curvature parameter c to 1;
- Lower the run-in age to 100;
- Change the start age in the graduation from 60 to 70 and 80; and
- Use existing ONS exposure estimates (as opposed to revised estimates from Section 2).

To illustrate the sensitivity to the choices we have made regarding graduation, we show the impact on period life expectancy for alternative choices of the key graduation parameters for males and females in Tables 4.2 and 4.3 below. We have highlighted the parameter that has been changed from the base case in **bold light blue text**. Note that for the sensitivities where we have limited the graduation age to below 105, we have reduced the run-in age accordingly.

Table 4.2: Female Sensitivities

Data	New KT	New KT	New KT	New KT	New KT	New KT	New KT	New KT	New KT	Old KT
Start Age	60	60	60	60	60	60	60	**70**	**80**	60
End Age	113	113	**100**	**105**	**109**	113	113	113	113	113
G(x)	5	**2****	5	5	5	5	5	5	5	5
Run-in age	105	105	**100**	105	105	105	**100**	105	105	105
c *	1.25	1.25	1.25	1.25	1.25	**1**	1.25	1.25	1.25	1.25
Life Expectancy										
65	20.26	20.18	20.26	20.26	20.26	20.26	20.26	-	-	20.25
75	12.54	12.52	12.54	12.54	12.54	12.54	12.54	12.54	-	12.53
85	6.61	6.69	6.61	6.61	6.61	6.61	6.61	6.61	6.62	6.60
95	3.10	3.03	3.10	3.10	3.10	3.10	3.10	3.10	3.09	3.12
% Change										
65	-	-0.39%	0.00%	0.00%	0.00%	0.00%	0.00%	-	-	-0.05%
75	-	-0.16%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	-	-0.08%
85	-	+1.21%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	+0.15%	-0.15%
95	-	-2.26%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	-0.32%	+0.65%

* where c is the curvature parameter.

** Given the G(7) graduation had the best information criteria, we have also considered this sensitivity. For females it produces equivalent life expectancies to the base case at all ages other than 95 where it is 0.2% lower.

From Table 4.2 we can see that, as expected, the G(2) graduation produces lower life expectancies at age 95 although interestingly higher life expectancies at age 85 due to the G(2) curve being lower at younger ages. The rest of the sensitivities, including raising the start age of graduation, have very little impact. The sensitivity with respect to the ONS KT method of estimating exposures leads to a slightly higher life expectancy at age 95 and very marginally lower at younger ages.

Table 4.3: Male Sensitivities

Data	New KT	New KT	New KT	New KT	New KT	New KT	New KT	New KT	Old KT
Start Age	60	60	60	60	60	60	**70**	**80**	60
End Age	109	109	**100**	**105**	109	109	109	109	109
G(x)	5	**2****	5	5	5	5	5	5	5
Run-in age	105	105	**100**	105	105	**100**	105	105	105
c	1.25	1.25	1.25	1.25	**1**	1.25	1.25	1.25	1.25
Life Expectancy									
65	17.56	17.51	17.56	17.56	17.56	17.56	-	-	17.57
75	10.62	10.63	10.62	10.62	10.62	10.62	10.62	-	10.63
85	5.58	5.63	5.58	5.58	5.58	5.58	5.59	5.60	5.61
95	2.71	2.60	2.70	2.71	2.71	2.72	2.71	2.69	2.82
% Change									
65	-	-0.28%	0.00%	0.00%	0.00%	0.00%	-	-	+0.06%
75	-	+0.09%	0.00%	0.00%	0.00%	0.00%	0.00%	-	+0.09%
85	-	+0.90%	0.00%	0.00%	0.00%	0.00%	+0.18%	+0.36%	+0.54%
95	-	-4.06%	-0.37%	0.00%	0.00%	+0.37%	0.00%	-0.74%	+4.06%

* where c is the curvature parameter.

** Given the G(7) graduation had the best information criteria we have also considered this sensitivity. For males it produces equivalent life expectancies to the base case at all ages other than 85 where it is 0.1% higher and 95 where it is 0.7% lower. However, this graduation plateaus around age 104 and so is not preferred due to the necessary unsmooth progression to a value of $\mu_{120} = 1$.

We see a similar pattern to females in the male sensitivities, albeit the reduction in life expectancy for moving to a G(2) graduation is larger. We also see that the use of the revised population exposures leads to a larger reduction in life expectancy, of around 4%, at age 95.

We therefore conclude that our graduation approach is reasonable.

Comparison with English Life Tables

A further interesting comparison of our graduated rates is to compare them against the English Life Tables as produced by the ONS. The high age treatment of these tables is summarised in Working Paper 85 and broadly involves using a weighted average of values produced by Gompertz and logistic models with a range of transition ages. To make the comparison we have applied the graduation as proposed above over the period 2010 – 2012 to be consistent with the data underlying the ELT 17 tables. Chart 4.11 and Table 4.4 set out how the rates compare.

Chart 4.11: Comparison of mortality rates for proposed graduation with ELT 17, 2010-2012 for males (LHS) and females (RHS)

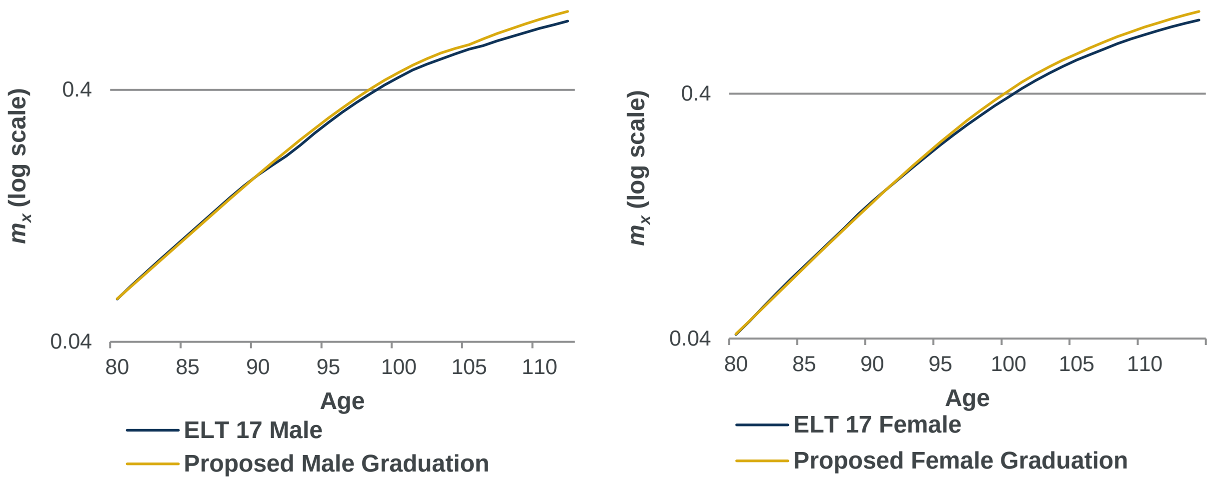

Table 4.4: Comparison of life expectancies for proposed graduation with ELT 17 life expectancies

Life Expectancy	Proposed Male 2010 - 2012	ELT 17 Male	Difference	Proposed Female 2010 - 2012	ELT 17 Female	Difference
65	18.32	18.32	0.00%	20.88	20.88	0.00%
75	11.17	11.19	0.18%	13.01	13.03	0.15%
85	5.81	5.84	0.52%	6.85	6.85	0.00%
95	2.75	2.84	3.27%	3.16	3.26	3.16%

These results demonstrate that there is broad consistency between the two approaches up to age 85 but that the mortality rates produced using this example of the proposed methodology are higher than for ELT17 at higher ages. This is consistent with our findings on extinct generation mortality in Working Paper 85.

4.3 Case studies

The framework described in Section 4.1 has been applied to the “08” Series annuitant mortality and “S2” Series tables to illustrate the results of the approach. The extensions applied are relative to the graduations of ONS data described in Section 4.2.

Case study 1: 08 Series

We illustrate the framework in some detail for PML08 (pension annuities in payment, male, lives), before presenting the results for other “08” Series annuities tables.

PML08 extension

Step 1: Graduate portfolio data

Chart 4.12 shows the published graduation of the PML08 data (before high age extension) relative to our graduation of ONS data for males.

Chart 4.12: Comparison of PML08 data with our graduation of ONS data, males

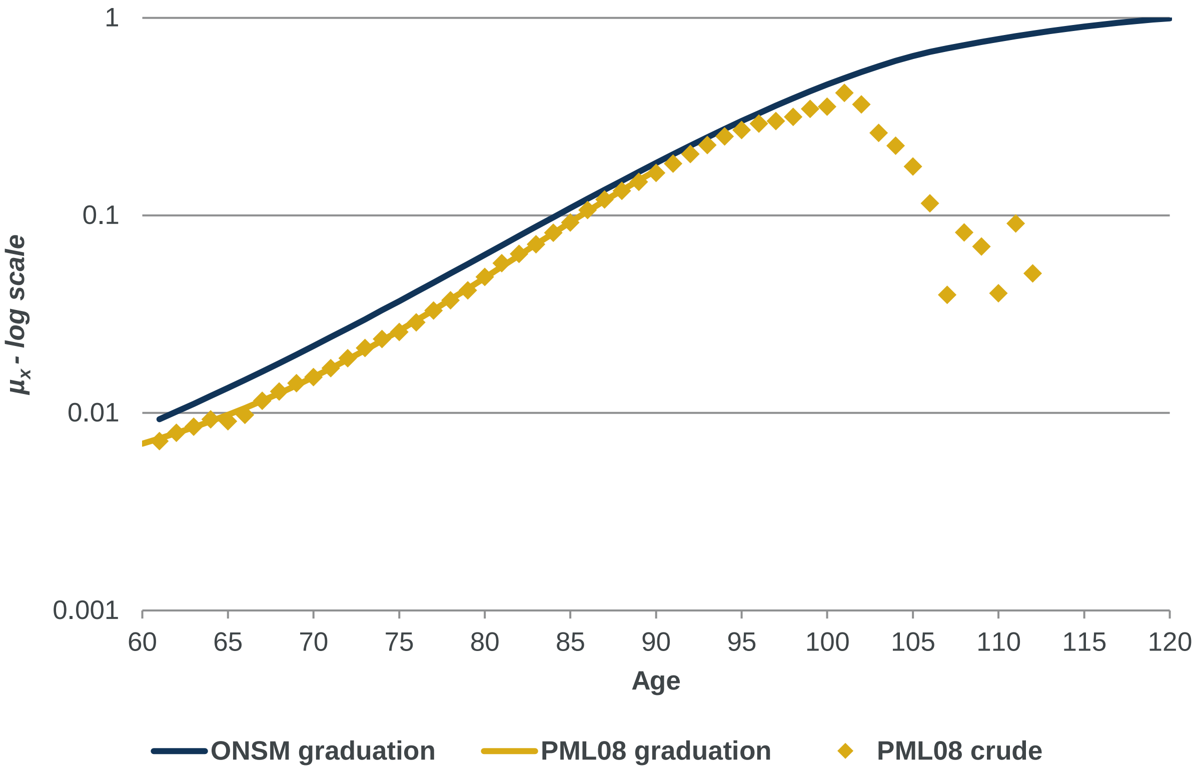

The graduated values were used to age 90 (gold line). The crude mortality rates (gold markers) are increasingly volatile above age 90, and particularly above age 100.

Step 2: Analyse convergence

The convergence between the PML08 dataset and our graduation of ONS data for males is more apparent when presented as a ratio, as in Chart 4.13. This shows a comparison of the crude mortality rates underlying the PML08 table with our graduation of ONS data for males. A pattern of convergence is seen from around age 75 up to around age 95.

Chart 4.13: Convergence between PML08 dataset and our graduation of ONS data, males

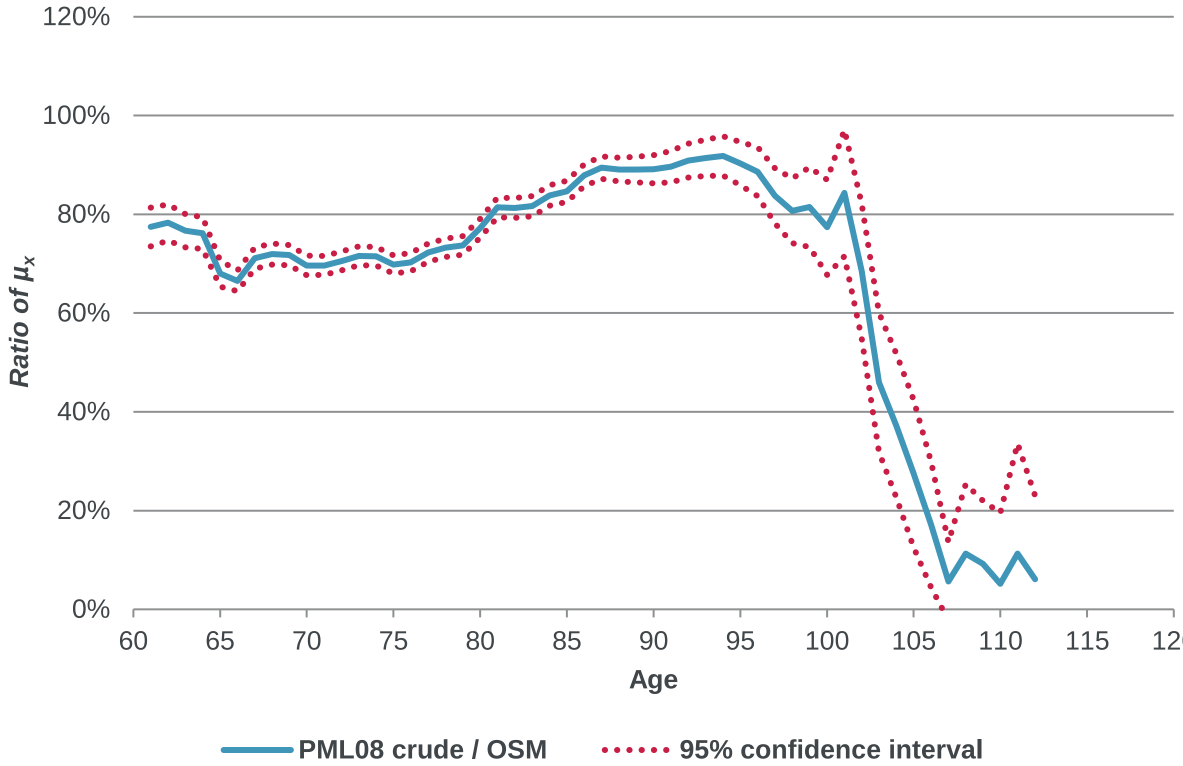

The rate of convergence above age 75 between portfolio and population data averages 33% every 5 years over the ages 75 to 90. The rate of convergence is weighted to ages with most data – meaning that very little weight is given at ages above 95 where there is less data and it may be less reliable Between age 75 and the high 80s there is a clear pattern of convergence (with the ratio shown increasing rapidly towards 100%). From the high 80s to age 95 the gradient is lower, although there is evidence of continued convergence over each 5-year age range to age 95.

Looking at convergence patterns for other “08” Series annuities tables there is some volatility in the rates of convergence at different ages in each case. In some cases, the observed rate of convergence is lower at the top of the range analysed (as per PML08, shown above); in other cases, the rate of convergence is higher at high ages. In view of the volatility across the age range (and widening confidence intervals at higher ages) we focus on the average convergence rate across the age range for which a broadly similar pattern is observed, rather than seeking to draw conclusions that convergence is becoming more or less rapid at the highest ages.

In this example, convergence has been analysed from age 75. Below age 75 the differential between the population table and portfolio mortality rates remains quite stable (and so is a different pattern to that seen in the run up to and around age 90, the age from which we wish to extend the graduation). In general, we start the convergence analysis from the youngest age from which a broadly consistent pattern emerges (age 75 in this case). It is acknowledged that the age from which a broadly consistent pattern of convergence is seen is not objectively defined, and does require expert judgement in its selection.

Broadly similar rates of convergence were observed for most other “08” Series annuities tables, providing some confidence that it may be reasonable to extend the graduation using a projected rate of convergence of that order for each of those datasets.

Step 3: Extend graduation

Chart 4.14 shows the table extended above age 91 assuming a continued rate of convergence (to our graduation of ONS data for males) of 30% every 5 years.

Chart 4.14: Comparison of PML08, extended above age 91 using our methodology with our graduation of ONS data, males

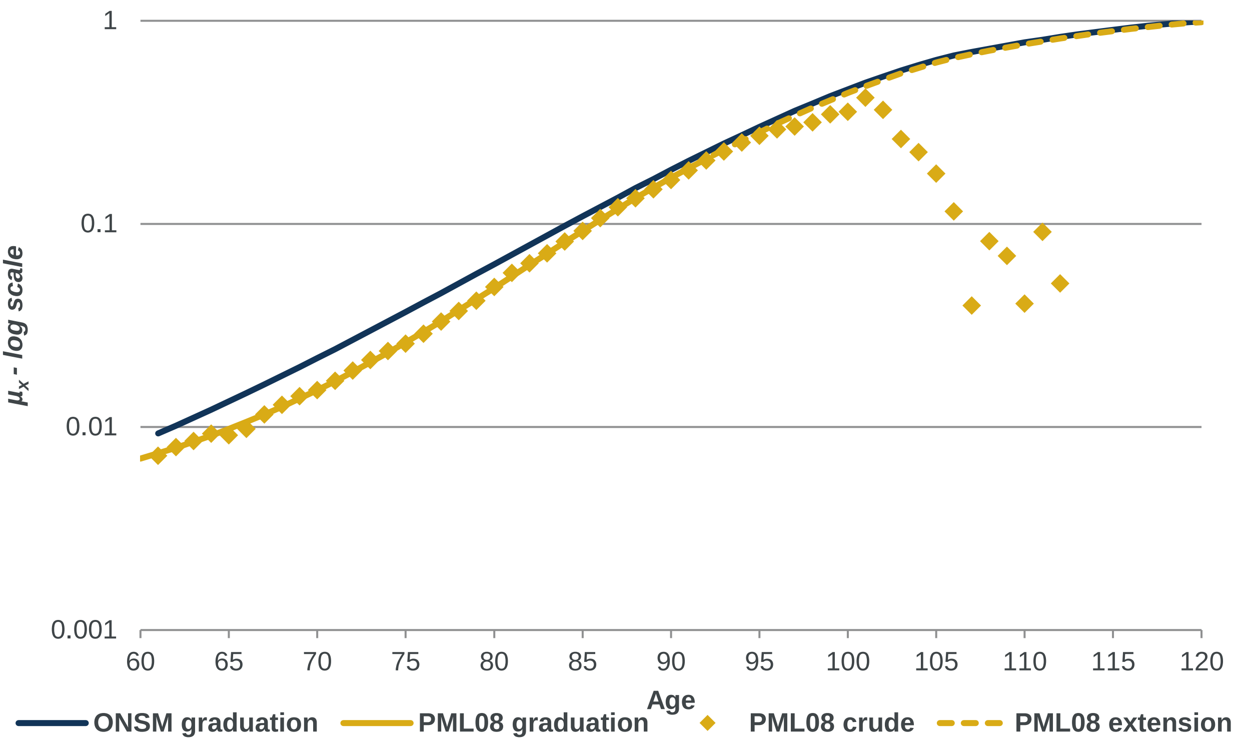

It is interesting to compare this alternative extension to that in the published PML08 table. This is best seen by considering ratios, as shown in Chart 4.15.

Chart 4.15: Comparison of PML08 with our alternative extension method and ONS data, males

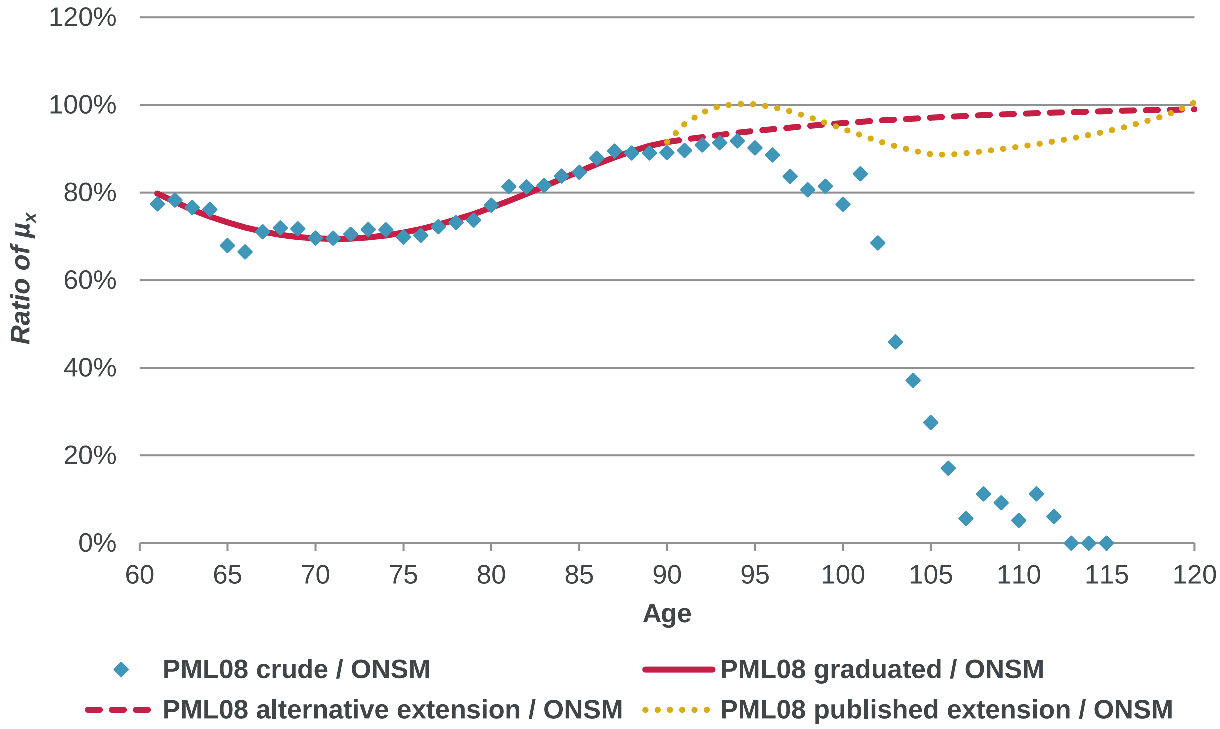

It can be seen that the published "08" Series extensions initially converge rapidly towards population data (i.e. the 100% line), before diverging to lower assumed rates of mortality at the highest ages. Our alternative extension (by construction) converges in a steadier fashion towards our graduation of ONS data for males. It also appears a better fit to the underlying PML08 crude rates between ages 90 and 95.

Our choice of extension approach (specifically assuming a continued rate of convergence towards our population graduation) is one of many possible alternative formulations. In choosing this approach we have given some consideration to alternatives and provide a few comments below on our reasoning for taking the approach we have, rather than another relatively simple alternative (specifically to assume log linear convergence to the population rates).

- Having considered the pattern of convergence between various populations, we have observed that in many cases the populations tend towards each other at high ages, but data volumes become too small to conclude that they have actually converged fully by any given age. For example, considering charts such as Chart 4.13, we see cases which look more like rates tending towards each other, than strong evidence of convergence by a specific age.
- Where possible, we wish the extension to be consistent with underlying metrics that we have observed for the dataset in question, or for other datasets which we may believe to be relevant. The approach we have taken to analysing convergence for the age range for which we have data, provides a simple metric which can be used to define the extension.
- Assuming a log linear convergence to the population rates requires selection of an age by which full convergence has occurred. We find it difficult to hypothesise by what age full convergence will have occurred from the analyses we have undertaken.

Table 4.5 illustrates the impact of the alternative extension on life expectancy at various ages.

Table 4.5: Impact of alternative extension methodology on life expectancy for PML08

Age	Period life expectancy* (published table)	Period life expectancy (alternative extension)	Impact
65	19.343	19.378	+0.18%
75	11.730	11.772	+0.36%
85	5.908	5.978	+1.18%
95	2.765	2.829	+2.31%

*Note that these figures differ slightly to those published, as the work undertaken for this paper used mortality rates at integer ages only hence a more approximate conversion from μ_x to q_x has been necessary, leading to slight differences in life expectancy statistics.

Other tables

This process has been repeated for PFL08, PFA08, PMA08, LFL08 and LML08. Table 4.6 summarises the ages from which extensions are applied (which have been inherited from the "08" Series tables in each case) and the observed rate of convergence to our graduations of ONS data for males or females as appropriate at higher ages. In general, we would suggest that it is preferable to use graduated portfolio data to as high an age as the data can support, i.e. until data volumes fall away or until there are concerns about the accuracy of data. Such an approach limits the impact of the choice of extension and has regard to the portfolio data over the age range where it is thought to be reliable and relevant. Considering each graduation in isolation in this way may lead to extensions being applied from different ages; it is therefore necessary to balance the desire to maximise use of data and the potential downside of extensions being less consistent if applied from different ages.

Table 4.6: Summary of high age extensions and convergence for “08” Series tables

Table	Age from which high age extension is applied	Age above which consistent convergence pattern is observed	Average convergence rate (per 5years) from that age
PFL08	91	75	30%
PFA08	91	75	21%
LFL08	90	75	60%
PML08	90	75	33%
PMA08	90	75	28%
LML08	90	75	43%

Based on the above, all tables have been extended using an assumed rate of convergence of 30% every 5 years (even though the rate of convergence appears higher for the life annuities tables, LFL08 and LML08).

Charts 4.16 and 4.17 show the resulting extensions for males and females respectively.

Chart 4.16: Extensions for male “08” Series tables relative to our graduation of ONS male data

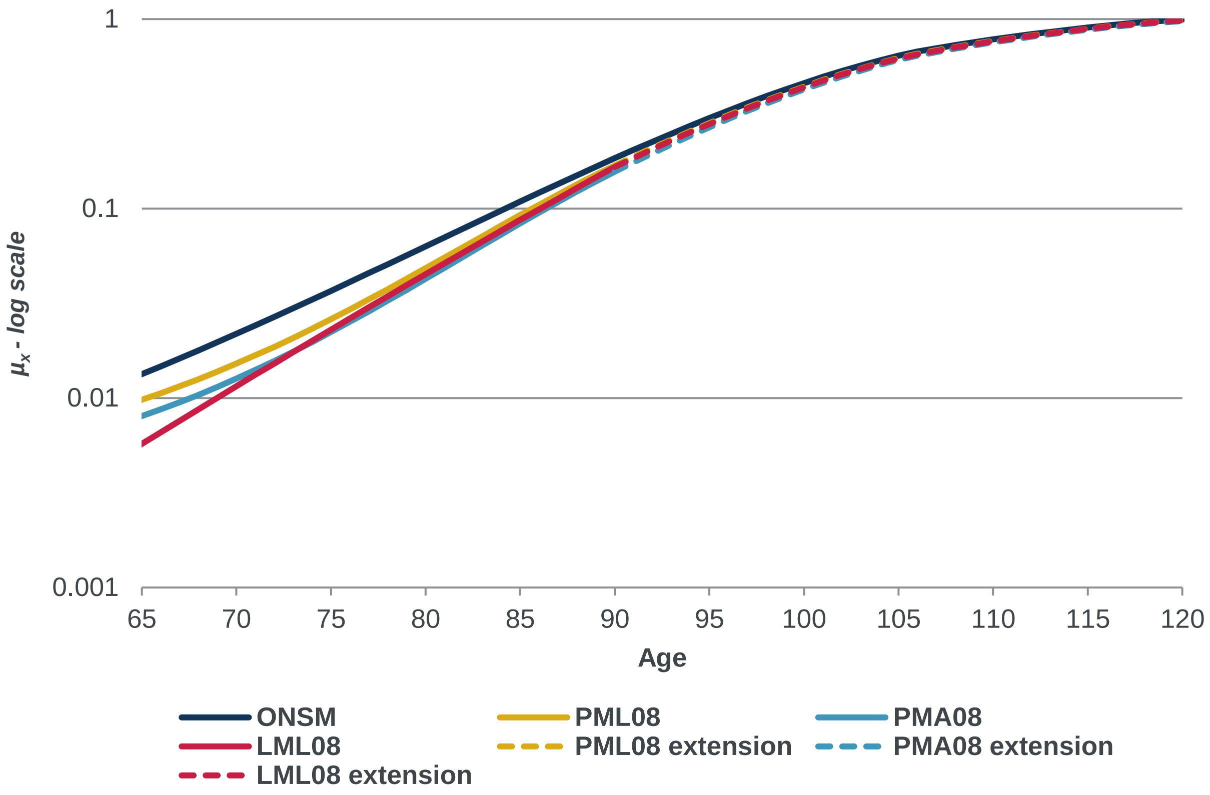

Chart 4.17: Extensions to female "08" Series tables relative to our graduation of ONS female data

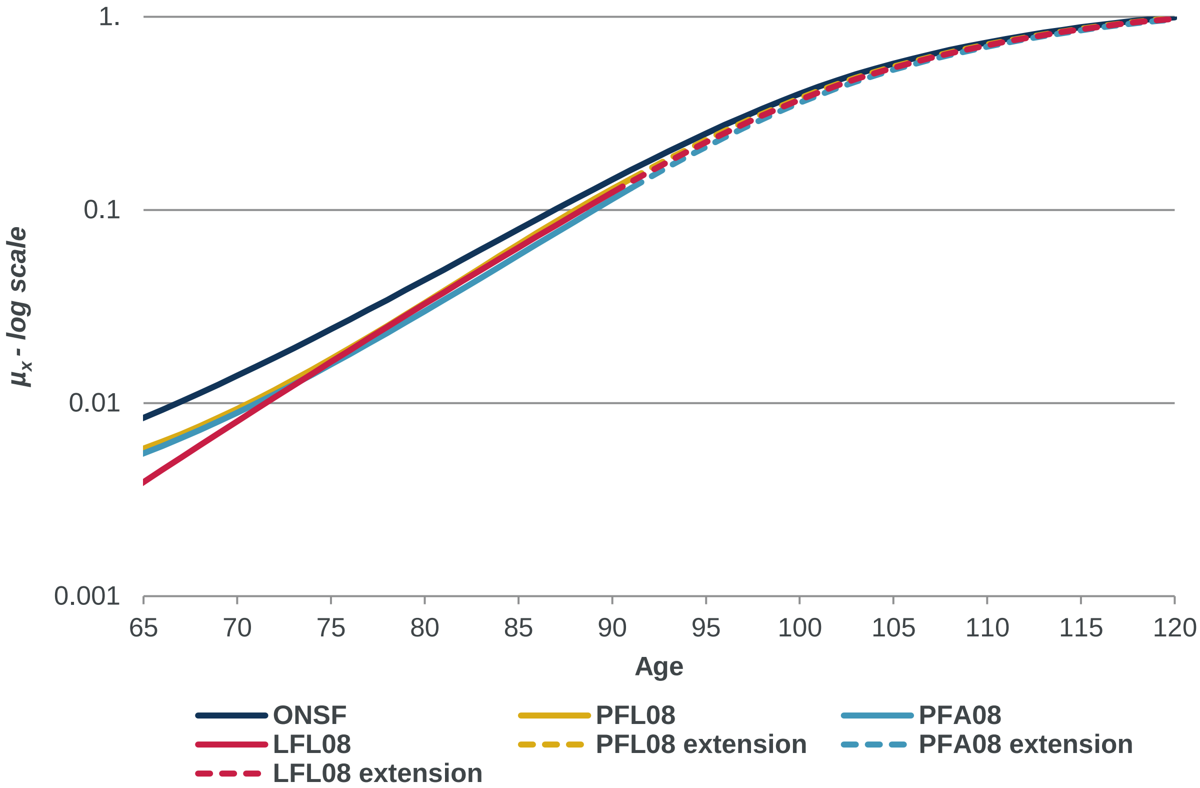

It can be seen that this extension approach, combined with the same rate of convergence for the various tables, ensures that there is no crossing of extensions.

Charts 4.18 and 4.19 compare these extensions (dashed lines) with those published in the "08" Series tables (dotted lines). In each case they are displayed relative to our graduations of ONS data for males or females as appropriate.

Chart 4.18: Comparison of actual "08" Series tables with our alternative extensions, males

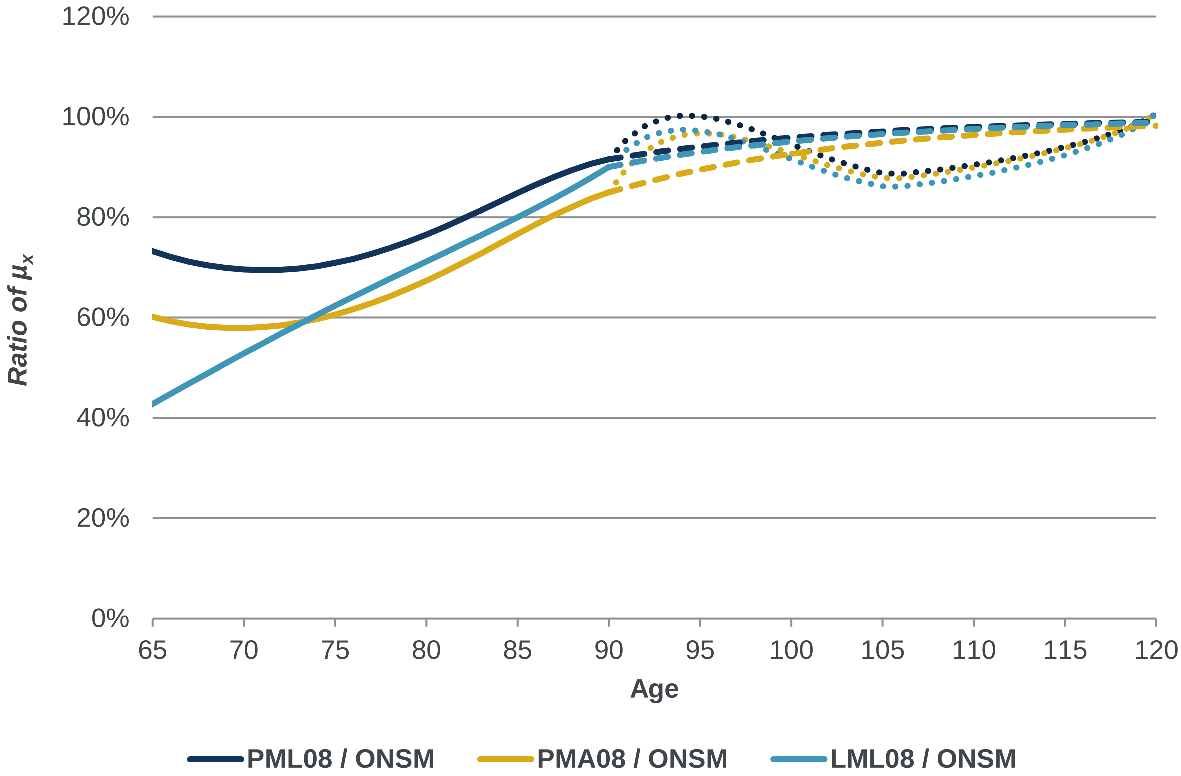

Chart 4.19: Comparison of actual "08" Series tables with our alternative extensions, females

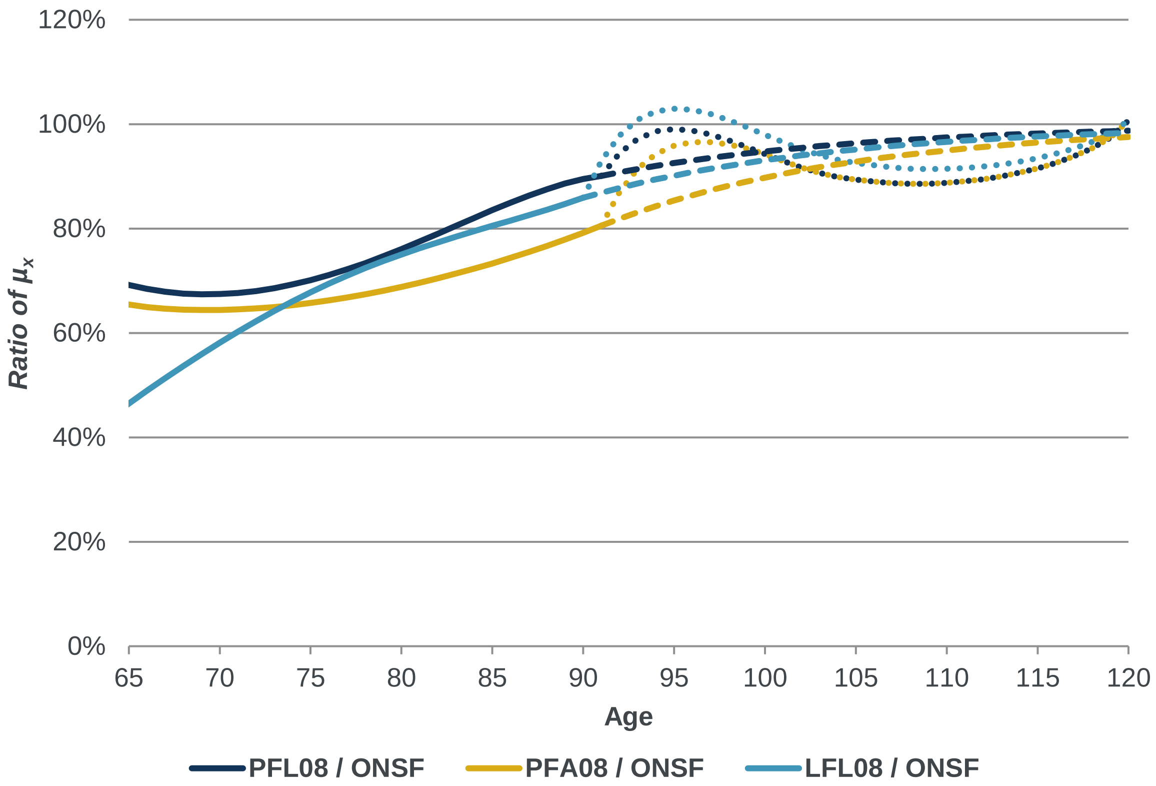

As for PML08, it can be seen that the published "08" Series extensions initially converge rapidly towards population data (i.e. the 100% line), before diverging to lower assumed rates of mortality at higher ages. These alternative extensions (by construction) converge in a steadier fashion towards graduated population rates.

Table 4.7 summarises the impact on period life expectancy of these extensions. It can be seen that life expectancy increases in all cases. The impact is higher for females (relative to males) and at the highest ages.

Table 4.7: Impact of alternative extension methodology on period life expectancy for selected "08" Series tables

Table	Period life expectancy* (published table)				Period life expectancy (alternative extension)				Impact			
	65	75	85	95	65	75	85	95	65	75	85	95
PFL08	21.94	13.64	7.02	3.18	21.99	13.69	7.10	3.26	+0.23%	+0.37%	+1.14%	+2.52%
PFA08	22.55	14.22	7.43	3.22	22.65	14.33	7.58	3.43	+0.44%	+0.77%	+2.02%	+6.52%
LFL08	22.24	13.72	7.05	3.08	22.38	13.87	7.25	3.32	+0.63%	+1.09%	+2.84%	+7.79%
PML08	19.34	11.73	5.91	2.77	19.38	11.77	5.98	2.83	+0.21%	+0.34%	+1.18%	+2.17%
PMA08	20.23	12.32	6.19	2.83	20.28	12.38	6.28	2.93	+0.25%	+0.49%	+1.45%	+3.53%
LML08	20.19	12.08	6.05	2.84	20.22	12.11	6.11	2.85	+0.15%	+0.25%	+0.99%	+0.35%

*Note that these figures differ slightly to those published, as the work undertaken for this paper used mortality rates at integer ages only hence a more approximate conversion from μ_x to q_x has been necessary, leading to slight differences in life expectancy statistics.

Case Study 2 – SAPS pensioner data

A similar process has been followed in relation to "S2" Series tables.

S2PFL, S2PFA, S2PML, S2PMA, S2DFL, S2DFA, S2NFA, S2NMA, S2IFA and S2IMA extensions

This section illustrates the results other than for light, medium and heavy tables.

In each case extensions are applied from age 95 (inherited from the "S2" Series). The age above which a consistent pattern (of convergence or otherwise) occurs is more variable for the "S2" Series tables than it was for the "08" Series tables, as illustrated in Table 4.8.

Table 4.8: Summary of high age extensions and convergence for selected "S2" Series tables

Table	Age from which high age extension is applied	Age above which consistent convergence (or non-convergence) pattern is observed	Average convergence rate (per 5years) from that age
S2PFL	95	80	-4%
S2PFA	95	70	24%
S2NFA	95	70	27%
S2IFA	95	70	35%
S2DFL	95	80	-32%
S2DFA	95	70	7%
S2PML	95	60	29%
S2PMA	95	60	20%
S2NMA	95	60	22%
S2IMA	95	65	24%

With the exception of the dependants and female lives tables, the observed rate of convergence is typically around 20-30% every 5 years. Given the variation implicit in crude mortality rates, the above metrics are also subject to random error. We also note that using the same convergence rate to project different tables will (when extending from the same age and relative to the same population table) avoid crossings in the extensions. Taking a pragmatic approach, we have extended these tables using an assumed rate of convergence of 25% each 5 years.

Table 4.8 highlights that for some datasets the convergence metric can be close to zero (e.g. S2DFA) implying a low rate of convergence or even divergence (e.g. S2PFL and S2DFL). In such cases inspecting the underlying convergence patterns (rather than relying on a single convergence metric) can provide a better understanding of what the underlying data is showing and inform the decision on the extension.

In the case of S2PFL, there is little convergence observed between 80 and 95, and some strong divergence (albeit with large confidence intervals, and potentially based on suspect data) above age 95. Given the divergence appears to be driven by potentially unreliable high age data, it would seem unwise to assume divergence in the extension. It is then a question of whether it is preferable to take a pragmatic approach of assuming convergence (from 93% of the population table at age 95) in the extension, or follow the trend the data seems to suggest at ages 80-95 and assume a convergence rate of nil%. We have illustrated the former approach, which is consistent with the approach taken for S2PFA (which showed strong evidence of convergence) and so avoids introducing inconsistencies between the tables in the high ages extensions.

Similar considerations can be taken in forming a view on S2DFA and S2DFL.

Charts 4.20 and 4.21 compare these extensions (dashed lines) with those published in the “S2” Series tables (dotted lines). In each case they are displayed relative to our graduations of ONS data for males or females as appropriate.

Chart 4.20: Comparison of actual “S2” Series tables with our alternative extensions, males

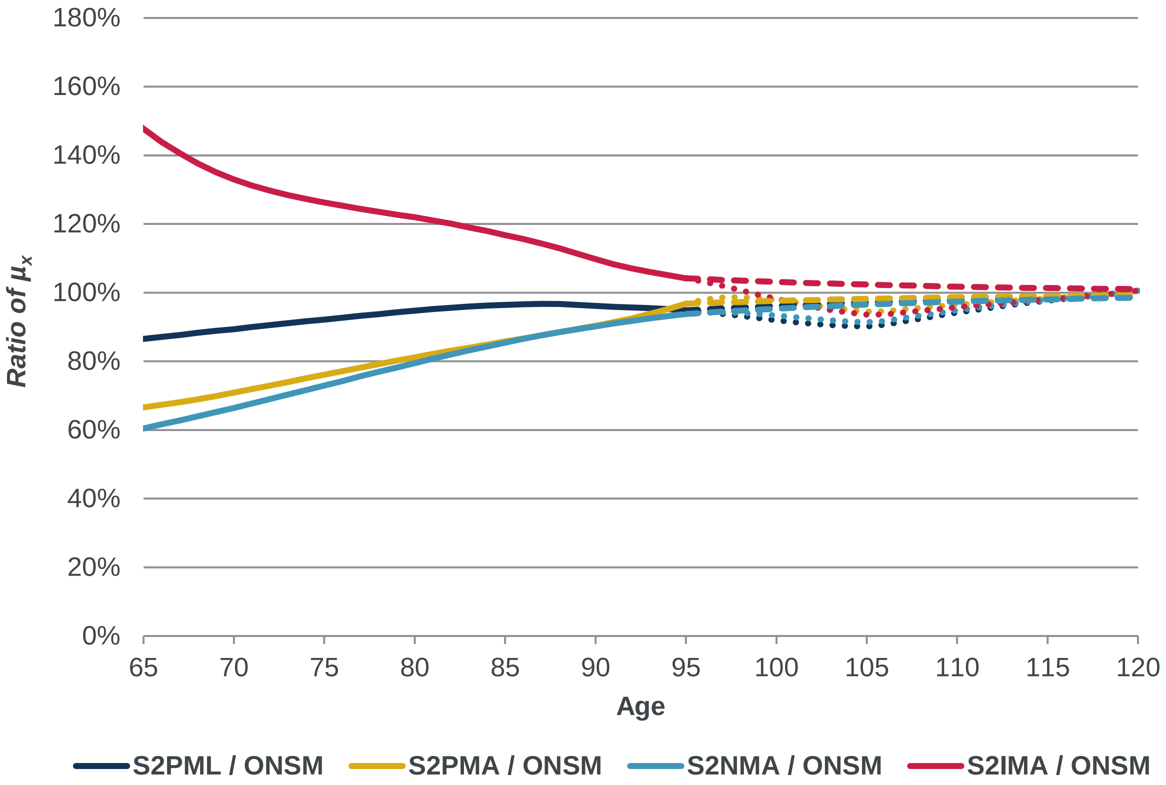

For males, the published “S2” series extensions diverge from our population graduation (i.e. the 100% line) before converging again at higher ages. In contrast these alternative extensions (by construction) converge in a steady fashion towards graduated population rates.

It can be seen that this extension approach, combined with the selected rates of convergence does not result in crossing of extensions.

Chart 4.21: Comparison of actual "S2" Series tables with our alternative extensions, females

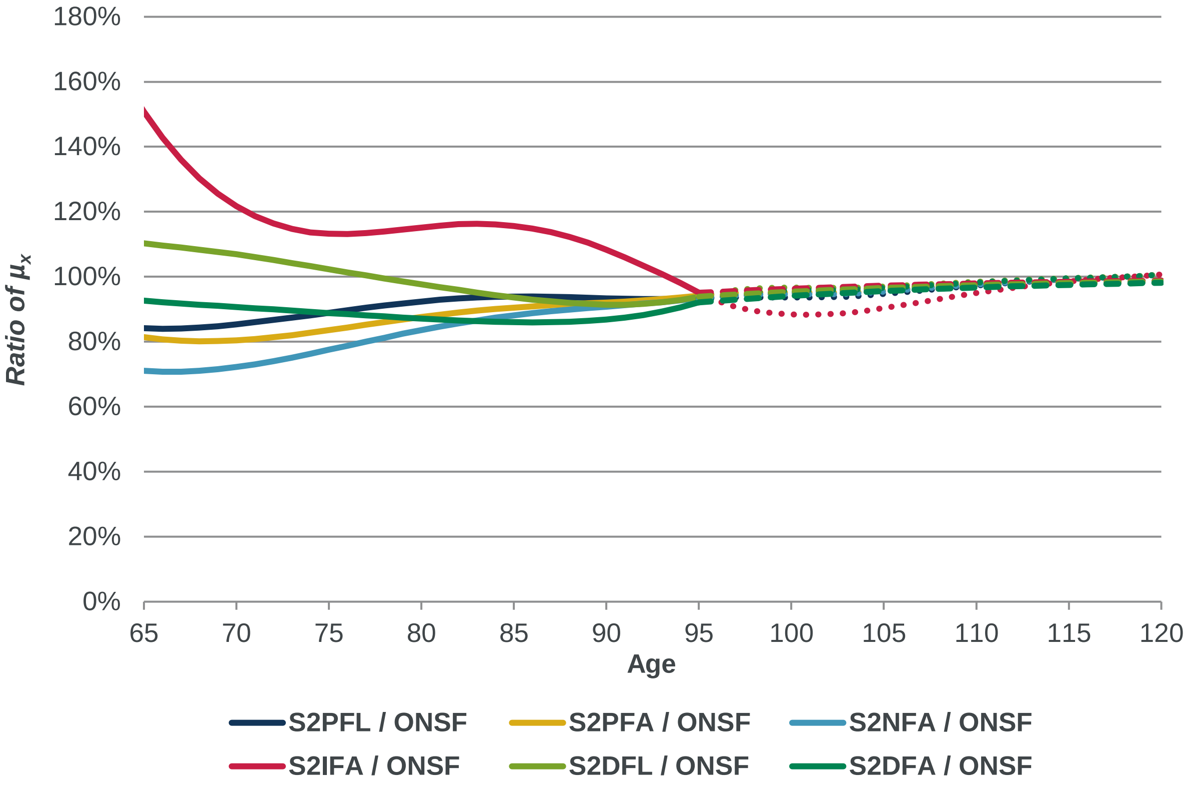

For females (excluding ill-health), the published "S2" Series extensions converge relatively smoothly towards population data (i.e. the 100% line). These alternative extensions (by construction) also converge in a steady fashion towards graduated population rates.

It can be seen that this extension approach, combined with the selected rates of convergence does not result in crossing of extensions.

For S2IFA, the alternative extension converges more rapidly to population rates than the "S2" Series extension. This results in a "kink" in the curve at age 95 (i.e. a step change in the gradient), but avoids the graduation crossing other tables in the extension. It is a matter of judgement as to which feature is preferred, although inspection of the underlying data does not provide particular support for the gradient of S2IFA at age 95, hence the different gradient in the extension does not seem to conflict with the underlying data.

Table 4.9 summarises the impact on period life expectancy of these extensions. Generally, for the "S2" Series tables shown in Table 4.9 the change in life expectancy would be very small, although some larger impacts are observed at very high ages for a few tables, notably the ill-health tables.

Table 4.9: Impact of alternative extension methodology on period life expectancy for selected "S2" Series tables

Table	Period life expectancy* (published table)				Period life expectancy (alternative extension)				Impact			
	65	75	85	95	65	75	85	95	65	75	85	95
S2PFL	20.943	12.979	6.885	3.266	20.940	12.976	6.881	3.247	-0.01%	-0.02%	-0.06%	-0.58%
S2PFA	21.227	13.181	6.941	3.231	21.227	13.180	6.941	3.229	+0.00%	+0.01%	-0.00%	-0.06%
S2NFA	21.582	13.371	7.006	3.250	21.581	13.371	7.006	3.248	+0.00%	+0.00%	-0.00%	-0.06%
S2IFA	19.200	11.883	6.329	3.346	19.186	11.866	6.301	3.208	-0.07%	-0.14%	-0.44%	-4.12%
S2DFL	20.349	12.812	6.924	3.208	20.353	12.816	6.931	3.236	+0.02%	+0.03%	+0.10%	+0.87%
S2DFA	21.144	13.326	7.139	3.242	21.150	13.332	7.148	3.277	+0.03%	+0.05%	+0.13%	+1.08%
S2PML	18.073	10.902	5.734	2.866	18.070	10.898	5.726	2.816	-0.02%	-0.04%	-0.14%	-1.74%
S2PMA	19.194	11.593	5.977	2.759	19.195	11.595	5.979	2.776	+0.01%	+0.02%	+0.03%	+0.62%
S2NMA	19.412	11.686	6.009	2.856	19.410	11.684	6.006	2.839	-0.01%	-0.02%	-0.05%	-0.60%
S2IMA	15.933	9.656	5.182	2.680	15.931	9.653	5.176	2.632	-0.01%	-0.03%	-0.12%	-1.79%

*Note that these figures differ slightly to those published, as the work undertaken for this paper used mortality rates at integer ages only hence a more approximate conversion from μ_x to q_x has been necessary, leading to slight differences in life expectancy statistics.

"S2" Series light, middle and heavy extensions

Extensions for the "S2" Series light, medium and heavy tables have been extended with reference to their respective (extended) "parent" tables. The light, middle and heavy tables are subsets of the SAPS dataset, based on different pension amounts, where lives with higher pensions are observed to experience lighter mortality.

In each case extensions are applied from age 90 (inherited from the "S2" Series). This is younger than for other "S2" Series tables, which were extended from age 95. It might be the case that it is necessary to extend child tables from a younger age than parent tables as the data volumes reduce from a younger age for these smaller datasets; however, this will not necessarily be the case for all child graduations and in general the data should be graduated to the highest possible age.

As for other "S2" Series tables, the age above which a consistent pattern (of convergence or otherwise) to the parent table occurs is variable, as illustrated in Table 4.10. Convergence for child tables should be analysed relative to their respective (extended) parent tables, rather than with the population table, if the parent table is viewed to be more directly relevant than the population table (which would usually be the case).

Table 4.10: Summary of high age extensions and convergence with parent tables for “S2” Series light, middle and heavy tables

Table	Age from which high age extension is applied	Age above which convergence is observed	Average convergence rate (per 5 years) from that age
S2PFA_L	90	70	54%
S2PFA_H	90	75	31%
S2PMA_L	90	75	36%
S2PMA_M	90	60	26%
S2PMA_H	90	75	33%
S2NFA_H	90	80	55%
S2NMA_L	90	75	35%
S2NMA_H	90	70	32%

While there is some variation in the observed rate of convergence, it is typically around 30-50% every 5 years, i.e. typically more rapid than the convergence observed between the parent and population tables.

Based on the above, these tables have been extended using an assumed rate of convergence of 40% every 5 years.

Table 4.11 summarises the impact on period life expectancy of these extensions. The impact of adopting these alternative extensions on life expectancies is very small for the “S2” Series tables considered in Table 4.11, except at the highest ages.

Table 4.11: Impact of alternative extension methodology on period life expectancy for "S2" Series light, middle and heavy tables

Table	Period life expectancy* (published table)				Period life expectancy (alternative extension)				Impact			
	65	75	85	95	65	75	85	95	65	75	85	95
S2PFA_L	21.543	13.375	6.976	3.226	21.547	13.379	6.982	3.234	+0.02%	+0.03%	+0.09%	+0.25%
S2PFA_H	20.604	12.666	6.730	3.241	20.582	12.642	6.691	3.168	-0.11%	-0.19%	-0.58%	-2.25%
S2PMA_L	20.454	12.400	6.246	2.757	20.476	12.425	6.285	2.847	+0.11%	+0.20%	+0.62%	+3.26%
S2PMA_M	18.325	11.099	5.793	2.762	18.316	11.088	5.772	2.725	-0.05%	-0.10%	-0.36%	-1.34%
S2PMA_H	17.243	10.361	5.522	2.765	17.225	10.338	5.477	2.662	-0.10%	-0.22%	-0.81%	-3.73%
S2NFA_H	20.886	12.824	6.808	3.259	20.867	12.803	6.775	3.197	-0.09%	-0.16%	-0.48%	-1.90%
S2NMA_L	20.609	12.487	6.316	2.854	20.632	12.514	6.357	2.926	+0.11%	+0.22%	+0.65%	+2.52%
S2NMA_H	17.504	10.471	5.565	2.862	17.484	10.446	5.517	2.731	-0.11%	-0.24%	-0.86%	-4.58%

*Note that these figures differ slightly to those published, as the work undertaken for this paper used mortality rates at integer ages only hence a more approximate conversion from μ_x to q_x has been necessary, leading to slight differences in life expectancy statistics.

5. Conclusions and next steps

This paper builds on the findings from Working Paper 85, which summarised published literature on high age mortality, and provided early analysis of data and modelling features that may have a particular impact at high ages.

The key findings from our analyses in this paper are:

- We have set out our proposed framework for modelling population exposures at high ages that seeks to address limitations of the KT approach currently used by the ONS. We have concluded that it would be appropriate to make an explicit allowance for mortality trend; to reduce the age at and above which the KT approach is deployed; to adopt a more sophisticated "Lexis triangle" approach to determine counts at the start of the year; and to allow for adjustments to smooth out abnormal population exposures. If these recommendations are introduced then we observe a total reduction in period life expectancy which is small at ages below 85, but is more material (around 1.5%-2.0% reduction) for ages 90 and above, although the impact can vary depending on gender and the time period being considered.
- We have further explored the case for mortality deceleration at very high ages, to determine a suitable shape for mortality. We have considered data and analyses performed on three international data sources and concluded that it is hard to justify a log-linear fit to the E&W data. Based on the papers we have reviewed, in addition to those considered in Working Paper 85, and analysis we have performed, we believe that a mortality curve which makes allowance for deceleration at advanced ages is supportable for period mortality at population level. We also believe that a value of μ_{120} of 1 is justifiable and a reasonable assumption currently.
- We have set out our proposed framework for closing-off mortality tables at high ages. We divide the age range for consideration into three segments which can be determined with reference to portfolio mortality data and population mortality data where sufficient, relevant data exists and then expert judgement, as set out in the previous paragraph, to determine a limiting value for mortality rates at very high ages. The framework has been designed to enable both portfolio and sub-portfolio graduations to be closed off using a consistent methodology. We have applied our framework to recent CMI datasets for Self-Administered Pension Schemes and insured annuitants. We intend to work with CMI mortality-related committees to support the implementation of this framework for future mortality tables.

In Working Paper 85 the Working Party set out its ambition to work with CMI committees on data issues and modelling features at very high ages. We have been seeking mortality data for large UK pension schemes with which to perform both quantitative and qualitative analyses with the intention of producing a further Working Paper setting out our findings.

The Working Party welcomes feedback on this paper and can be contacted at HighAgeMortality@cmilimited.co.uk.

References

1. Andreev, "A Method for Estimating Size of Population Aged 90 and over with Application to the 2000 U.S. Census Data" (2004): https://www.demographic-research.org/volumes/vol11/9/11-9.pdf
2. Cairns, Blake, Dowd and Kessler, "Phantoms never die: living with unreliable population data" (2015): http://www.macs.hw.ac.uk/~andrewc/papers/JRSS2016.pdf
3. CMI Working Paper 71, "Final "S2" Series of Mortality Tables" (2014): https://www.actuaries.org.uk/learn-and-develop/continuous-mortality-investigation/cmi-working-papers/saps/cmi-wp-71
4. CMI Working Paper 81, Final "08" Series annuities tables (2015): https://www.actuaries.org.uk/learn-and-develop/continuous-mortality-investigation/cmi-working-papers/annuities/cmi-wp-81
5. CMI Working Paper 83, "Recent mortality in England & Wales" (2015): https://www.actuaries.org.uk/learn-and-develop/continuous-mortality-investigation/cmi-working-papers/mortality-projections/cmi-wp-83
6. CMI Working Paper 84, "The CMI Mortality Projections Model, CMI_2015" (2015): https://www.actuaries.org.uk/learn-and-develop/continuous-mortality-investigation/cmi-working-papers/mortality-projections/cmi-wp-84
7. CMI Working Paper 85, "Initial report on the features of high age mortality" (2015): https://www.actuaries.org.uk/learn-and-develop/continuous-mortality-investigation/cmi-working-papers/mortality-projections/cmi-wp-85
8. CMI Working Paper 91, "CMI Mortality Projections Model consultation – technical paper" (2016): https://www.actuaries.org.uk/learn-and-develop/continuous-mortality-investigation/cmi-working-papers/mortality-projections/cmi-working-papers-90-91-and-93
9. CMI Working Paper 99, "CMI Mortality Projections Model: Software user guide" and accompanying model (2017): https://www.actuaries.org.uk/learn-and-develop/continuous-mortality-investigation/cmi-working-papers/mortality-projections/cmi-working-papers-97-98-and-99
10. Gavrilov and Gavrilova, "Mortality measurement at advanced ages: a study of the Social Security Administration Death Master File" (2011)
11. Gavrilov and Gavrilova, "Mortality of Supercentenarians: Does It Grow with Age?" (2015)
12. Lexis, "Introduction to the Theory of Population Statistics" (1875): http://gdz.sub.uni-goettingen.de/de/dms/load/toc/?PPN=PPN527522651
13. Ouellette and Bourbeau, "Measurement of Mortality among Centenarians in Canada" (2014)
14. Office for National Statistics, "Accuracy of official high-age population estimates, in England and Wales: an evaluation" (2016): https://www.ons.gov.uk/releases/accuracyofofficialhighagepopulationestimatesanevaluation
15. Office for National Statistics, "Quality and Methodology Information: Estimates of the Very Old (including Centenarians)" (2016): https://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/ageing/qmis/populationestimatesoftheveryoldincludingcentenariansqmi
16. Office for National Statistics, assorted death and population data publications for England & Wales:
    - Annual death and mid-year estimated population counts for 1961-2014 inclusive from "Population estimates for England and Wales 1961 to 2014 (Single year of age 0 to 105+)": https://www.ons.gov.uk/peoplepopulationandcommunity/populationandmigration/populationestimates/adhocs/005825populationestimatesforenglandandwales1961to2014singleyearofage0to105
    - Death counts for 2015 from "Death registrations by single year of age": https://www.ons.gov.uk/file?uri=/peoplepopulationandcommunity/birthsdeathsandmarriages/deat

hs/datasets/deathregistrationssummarytablesenglandandwalesdeathsbysingleyearofagetables/2015/previous/v1/syoa19742015publication.xls

- Mid-year estimated population counts for 2015 from
    i. England & Wales, ages 0-89: https://www.ons.gov.uk/file?uri=/peoplepopulationandcommunity/populationandmigration/populationestimates/datasets/populationestimatesforukenglandandwalesscotlandandnorthernireland/mid2015/ukmye2015.zip ("MYE2_population_by_sex_and_age_for_local_authorities_UK_2015")
    ii. England, ages 90-105+: https://www.ons.gov.uk/file?uri=/peoplepopulationandcommunity/birthsdeathsandmarriages/ageing/datasets/midyearpopulationestimatesoftheveryoldincludingcentenariansengland/current/e2015.xls
    iii. Wales, ages 90-105+: https://www.ons.gov.uk/file?uri=/peoplepopulationandcommunity/birthsdeathsandmarriages/ageing/datasets/midyearpopulationestimatesoftheveryoldincludingcentenarianswales/current/w2015.xls

17. Office for National Statistics, assorted birth data publications:
    - Quarterly births during 1838-1937 from historic archives of "Birth Statistics, England and Wales (Series FM1)":
    - Monthly births during 1938-1991 from "Live births registered by month of occurrence , England and Wales, 1938 to 1991" (2016): https://www.ons.gov.uk/file?uri=/peoplepopulationandcommunity/birthsdeathsandmarriages/livebirths/adhocs/005316livebirthsregisteredbymonthofoccurrenceenglandandwales1938to1991/livebirthsbymonthofocc1938to1991webv1tcm77432402.xls
    - Monthly births during 1992-2013 from "Characteristics of Birth 2" (2014) Table 2b: http://www.ons.gov.uk/ons/rel/vsob1/characteristics-of-birth-2--england-and-wales/2013/rft-characteristics-of-birth-2.xls
18. Rau et al, “Where is the level of the mortality plateau?” (2016)
19. Thatcher, Kannisto and Andreev, "The Survivor Ratio Method for Estimating Numbers at High Ages" (2002): http://www.demographic-research.org/volumes/vol6/1/6-1.pdf
20. Vandeschrick, "The Lexis diagram, a misnomer" (2001): http://www.demographic-research.org/Volumes/Vol4/3/4-3.pdf
21. Vincent, "La mortalité des vieillards" (1951)
22. Wilmoth, Andreev, Jdanov, and Glei, "Methods Protocol for the Human Mortality Database" (2007): http://www.mortality.org/Public/Docs/MethodsProtocol.pdf

Continuous Mortality Investigation Limited ("CMI") is registered in England & Wales
Company number: 8373631
Registered Office: 7th floor, Holborn Gate, 326-330 High Holborn, London, WC1V 7PP

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

Website: www.cmilimited.co.uk (redirects to www.actuaries.org.uk)

Continuous Mortality Investigation Limited is wholly owned by the Institute and Faculty of Actuaries.